

UNCLASSIFIED

AD NUMBER
ADB264839
NEW LIMITATION CHANGE
TO Approved for public release, distribution unlimited
FROM Distribution authorized to U.S. Gov't. agencies only; Proprietary Info.; Jan 2001. Other requests shall be referred to U.S. Army Medical Research and Materiel Command, 504 Scott Street, Fort Detrick, MD 21702-5012.
AUTHORITY
USAMRMC ltr, dtd 28 July 2003

THIS PAGE IS UNCLASSIFIED

AD _____

Award Number: DAMD17-99-1-9034

TITLE: Ultrasound Imaging Initiative

PRINCIPAL INVESTIGATOR: J. Frederick Cornhill, Ph.D.,
Geoffrey Lockwood, Ph.D., Raj Shekhar, Ph.D.
Christopher Hazard, M.S.

CONTRACTING ORGANIZATION: The Cleveland Clinic Foundation
Cleveland, Ohio 44195

REPORT DATE: January 2001

TYPE OF REPORT: Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

DISTRIBUTION STATEMENT: Distribution authorized to U.S. Government agencies only (proprietary information, Jan 01). Other requests for this document shall be referred to U.S. Army Medical Research and Materiel Command, 504 Scott Street, Fort Detrick, Maryland 21702-5012.

The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision unless so designated by other documentation.

20010403 039

NOTICE

USING GOVERNMENT DRAWINGS, SPECIFICATIONS, OR OTHER DATA INCLUDED IN THIS DOCUMENT FOR ANY PURPOSE OTHER THAN GOVERNMENT PROCUREMENT DOES NOT IN ANY WAY OBLIGATE THE U.S. GOVERNMENT. THE FACT THAT THE GOVERNMENT FORMULATED OR SUPPLIED THE DRAWINGS, SPECIFICATIONS, OR OTHER DATA DOES NOT LICENSE THE HOLDER OR ANY OTHER PERSON OR CORPORATION; OR CONVEY ANY RIGHTS OR PERMISSION TO MANUFACTURE, USE, OR SELL ANY PATENTED INVENTION THAT MAY RELATE TO THEM.

LIMITED RIGHTS LEGEND

Award Number: DAMD17-99-1-9034

Organization: The Cleveland Clinic Foundation

Location of Limited Rights Data (Pages):

Those portions of the technical data contained in this report marked as limited rights data shall not, without the written permission of the above contractor, be (a) released or disclosed outside the government, (b) used by the Government for manufacture or, in the case of computer software documentation, for preparing the same or similar computer software, or (c) used by a party other than the Government, except that the Government may release or disclose technical data to persons outside the Government, or permit the use of technical data by such persons, if (i) such release, disclosure, or use is necessary for emergency repair or overhaul or (ii) is a release or disclosure of technical data (other than detailed manufacturing or process data) to, or use of such data by, a foreign government that is in the interest of the Government and is required for evaluational or informational purposes, provided in either case that such release, disclosure or use is made subject to a prohibition that the person to whom the data is released or disclosed may not further use, release or disclose such data, and the contractor or subcontractor or subcontractor asserting the restriction is notified of such release, disclosure or use. This legend, together with the indications of the portions of this data which are subject to such limitations, shall be included on any reproduction hereof which includes any part of the portions subject to such limitations.

THIS TECHNICAL REPORT HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION.

REPORT DOCUMENTATION PAGEForm Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE January 2001	3. REPORT TYPE AND DATES COVERED Annual (21 Dec 99 - 20 Dec 00)	
4. TITLE AND SUBTITLE Ultrasound Imaging Initiative			5. FUNDING NUMBERS DAMD17-99-1-9034	
6. AUTHOR(S) J. Frederick Cornhill, Ph.D., Geoffrey Lockwood, Ph.D., Raj Shekhar, Ph.D., Christopher Hazard, M.S.				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Cleveland Clinic Foundation Cleveland, Ohio 44195 E-MAIL: cornhill@bme.ri.ccf.org			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Medical Research and Materiel Command Fort Detrick, Maryland 21702-5012			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES Report contains color graphics.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Distribution authorized to U.S. Government agencies only (proprietary information, Jan 01). Other requests for this document shall be referred to U.S. Army Medical Research and Materiel Command, 504 Scott Street, Fort Detrick, Maryland 21702-5012.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) This objective of this project is to build a real-time 3D ultrasound imaging system for combat casualty care. The high frame rate necessary for real-time 3D imaging is obtained using a synthetic aperture beamforming technique. The technique uses a fraction of the transmit pulses required by a conventional imaging system and permits very rapid image acquisition with no degradation of image quality. A beamformer capable of generating data for real-time 3D imaging has been implemented using a network of high speed digital signal processors. Optimized assembly code has been written for this system. A new multi-layer transducer array has been developed for improved electrical matching and noise performance. A prototype array has been incorporated into a compact probe head designed to rock the array for real-time 3D imaging. Besides the scanner development, an extensive set of software tools have been developed to provide easy and accurate analysis of the resulting 3D data. These tools allow real-time viewing of arbitrary two-dimensional planes through the data set and will allow quantitative assessment and display of the 3D anatomy.				
14. SUBJECT TERMS Ultrasound, 3-D Imaging, Image Processing, Synthetic Aperture				15. NUMBER OF PAGES 52
				16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT Unlimited	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

TABLE OF CONTENTS

I.	FRONT COVER	1
II.	SF-298	2
III.	TABLE OF CONTENTS	3
IV.	INTRODUCTION	4
V.	BODY	5
VI.	KEY RESEARCH ACCOMPLISHMENT.....	9
VII.	REPORTABLE OUTCOMES	9
VIII.	CONCLUSIONS	10
IX.	APPENDICES	11

INTRODUCTION

Providing accurate medical imaging at an aid station or remote field hospital is difficult. Portable ultrasound instrumentation can be designed for these applications, but the expertise required for an accurate diagnosis can only be obtained through years of training. Much of the information in an ultrasound examination is obtained by exploiting the real-time nature of the imaging modality. A successful diagnosis relies on the skill of the diagnostician to transform mentally dynamic two-dimensional (2D) images into the complex three-dimensional (3D) anatomy. Locating anatomical landmarks and moving the scan plane throughout the volume of interest are both critical components of this process. Without extensive training, it would be difficult for a medical corpsman to perform this procedure.

Researchers have started exploring tele-medicine in combination with 3D ultrasound data acquisition as a solution to this problem. This combination could potentially transfer the skill required to scan the patient and make the diagnosis from a medical corpsman to an imaging expert. Unfortunately, acquiring a 3D ultrasound data set is difficult. Modern 3D ultrasound systems are essentially conventional scanners modified to collect a series of 2D images. The images are later 'stacked' to represent the 3D anatomy. Although modern scanners are designed to collect 2D images in real-time (20 2D images/s), 3D image acquisition is slow. Slow image acquisition introduces the problem of how to align adjacent 2D images collected at different times. The patient and imaging probe can be immobilized to reduce movement of the anatomy between adjacent images. Cardiac and respiratory gating can also be applied. Even in a carefully controlled clinical setting, the resulting 3D data set is often badly distorted. If the patient, the anatomy, or the transducer moves during 3D image acquisition, the data must be discarded. Providing the care needed to obtain 'good' 3D data in the clinic is troublesome, on a battlefield it would be very difficult.

We are building an ultrasound imaging system that would avoid these problems. The two key components of our system are the following: 1) a high speed scanner that can collect 3D data 40 to 80 times faster than current 3D imaging approaches, and 2) a set of software tools for rapid image manipulation and analysis at a remote site. Real-time 3D image acquisition eliminates the need for patient immobilization, and cardiac and respiratory gating. A medical corpsman would simply place a small probe on the patient and position the sample volume by viewing a real-time two-dimensional image of the anatomy. Once the probe is correctly placed, a three-dimensional data set would be recorded in real-time (0.05 s for each 3D data set). Following data acquisition, the images would be transmitted to a central hospital for post-processing and analysis. An expert clinician could 're-scan' the patient by looking at multiple 2D planes through the data sets, or examine a computer reconstruction of the three-dimensional anatomy. The 3D data set could also be analyzed quantitatively to calculate dynamic changes in the anatomy. The entire imaging procedure, including subject preparation, would take only a few minutes.

BODY

A description of progress in each of the areas outlined in the statement of work is given below.

Phase III: Month 12 – 24

Scanner/Array Development:

- Beamformer - Complete software development for 64-channel system and verify performance.

PARTIALLY COMPLETED - The assembly language routines for all 64 channels have been written and independently verified. The board vendor had unforeseen manufacturing problems, having to do with the availability of certain specialized memory parts used in the system. As a result, we have only 17 of 33 digital signal processing boards available for testing. This precludes testing the complete 64-channel system. The full system is expected by February 2001. We have tested and verified the algorithms for a 24 channel system using simulated receive signals. We have also characterized and fully developed the synchronization electronics required for the 64-channel system. All testing to date indicates that the 64-channel system will function as designed. Calibration procedures are also being developed to ensure accurate performance. A conference proceeding publication, which describes the beamformer hardware and software, is included in Appendix 1.

- Front end electronics - Based on results of Phase II, design, fabricate, and test 64-channel front end.

PARTIALLY COMPLETED - Using a prototype array from Tetrad, a single channel of the front end was evaluated and found to be acceptable. The complete 64-channel front end has been designed, but fabrication is not yet completed.

- Transmit electronics - Based on results of Phase II, design, fabricate, and test multi-channel transmit boards.

COMPLETED - The decision was made to use a multi-layer transducer technology. A lab-bench multi-channel system has been designed and tested successfully. A PC board has been designed and sent out for fabrication. The pulser is capable of producing a monocycle from 0 to 400 volts peak-to-peak into a 10 to 50 ohm load. The pulse frequency can be varied from 1 to 10 MHz and the pulse repetition rate can be as high as 7 kHz. These specifications exceed the requirements for the system.

The following two tasks were assigned to Tetrad Corporation under subcontract. Four reports describing the progress made on the transducer array and probe head are included in Appendix 2 of this report.

- Transducer Array - Based on results of Phase II, design, fabricate, and test prototype array.

COMPLETED - Tetrad has designed and built a prototype array. A complete evaluation of this array is included in the Appendix 2.

- Probe Head - Assemble probe and deliver completed probe with drive circuitry to Cleveland Clinic.

PARTIALLY COMPLETED - A probe was assembled using a prototype array. The drive circuitry and electrical connections were evaluated. Details of the analysis are included in Appendix 2 of this report. A completed probe with an array meeting the specifications of this project is not yet completed.

- Develop software for scanner control.

PARTIALLY COMPLETED - Control of the scanner involves the integration of four major components: beamformer, pulser, transducer, and scan converting display. Each individual component has its own controls and overall system control involves synchronizing all of these systems. Using a specially developed synchronization board, we are able to ensure that the motion control signals from the transducer are used to simultaneously activate the beamformer and high voltage pulser. This is all done under the control of assembly language routines running on the beamforming hardware. The display has been designed to be data driven, so that minimal controls are needed. Software for controlling the display itself is fully developed. The major task remaining is the development of software which can control the motion of the transducer. This has been delayed until the actual rocking probe hardware is available.

User Interface & Image Processing Software:

- Develop clinical toolbox for detection of shrapnel

NOT COMPLETED - Unavailability of ultrasound images showing shrapnel or trauma in general has hindered the development of this toolbox. We are now planning to collect such images at the Walter Reed Hospital with the help of Drs. Gerald Moses and Kenneth Curley of the Telemedicine and Advanced Technology Research Center (TATRC). These images will be two-dimensional (2D), because most hospitals do not have the 3D acquisition capability. We plan to apply feature detection techniques being developed for the breast imaging toolbox to the trauma ultrasound images when they become available. In the future, when the use of 3D ultrasound imaging becomes widespread, our shrapnel detection technique can be extended to accommodate the third dimension.

- Develop clinical toolbox for cardiac imaging

PARTIALLY COMPLETED – Dynamic multi-planar reformatting (MPR), 3D image segmentation and 3D image registration are the three major engineering components we are developing for the cardiac imaging toolbox. The development of dynamic MPR and 3D image registration has been completed, while the development of 3D image segmentation is still ongoing. We expect to complete the development of 3D image segmentation in first quarter of 2001. Below we describe each engineering component individually.

Dynamic MPR. The dynamic MPR capability is a simple yet powerful mechanism to visualize a time varying 3D image by viewing any oblique or orthographic cross-section, replayed at the original heart rate (hence, dynamic). Although conceptually simple, the implementation of dynamic MPR is computationally demanding. To animate a view at 30 frames/sec, approximately 4 million trilinear interpolations per second (TRIPS) or 84 million floating-point operations (56 million multiplications and 28 million additions) per second are required. We have achieved interactive performance (>30 frames/sec) through the use 3D texture mapping hardware, resident on high-end graphics boards. The specific hardware we have used is rated to perform 142 million TRIPS. We are going to describe the dynamic MPR implementation in several publications.

3D Image Registration. The registration capability allows us to compare anatomy and physiology serially. The details of 3D image registration are described in the manuscript, "Mutual information-based rigid and nonrigid registration of ultrasound volumes." This manuscript (see Appendix 3) is currently under peer-review for publication in IEEE Transactions on Medical Imaging. We will present another paper, "Multifunction extension of simplex optimization method for mutual information based registration of ultrasound volumes," at the SPIE Medical Imaging Symposium in San Diego in February 2001. The abstract of this manuscript (see Appendix 4) is provided with this report.

3D Image Segmentation. We are extending our earlier published work to detect the endo- and epicardial surfaces of the left ventricle in 3D. It is a two-step semiautomatic approach, in which a surface template, placed in close proximity of the organ of interest by the user, is refined automatically thereafter under the influence of image gradients. This approach combines the strengths of automated and manual techniques, namely, the objectivity of automated techniques and the robustness of manual techniques.

We have created the necessary GUI to display a 3D image, place a 3D surface template inside the 3D image, and transform the shape globally as well as locally. The surface template can be rotated, translated and scaled comprising global transformation modes. In addition, any vertex with a user-controlled neighborhood can be pulled out or pushed in perpendicular to the surface (local morphing) by a user-defined amount. Shape changes are reflected in all views. The ongoing work focuses on: (1) creation of efficient data structure for 3D shapes, (2) development of effective automatic shape refinement algorithm, (3) segmentation of endo- and epicardial surfaces in all frames for shape tracking, and (4) validation.

- Develop clinical toolbox for breast imaging

PARTIALLY COMPLETED – As part of the breast imaging toolbox, we are developing automated tools to analyze lesions observed on 2D breast ultrasound images and to resolve cysts from non-cysts. The use of ultrasound is limited to these tasks in the current clinical practice. Moreover, these tasks are subjective since they are performed manually.

Our ongoing engineering development focuses on (1) segmentation of an observed lesion from the background and (2) separation of cysts from non-cysts. The test images, which number 18, represent a wide range of conditions and include 8 cysts and 10 non-cysts. We are investigating a set of optimal descriptors to discriminate between cysts and non-cysts. The available descriptors are tumor boundary fluctuation, tumor boundary roughness, geometric surface roughness, and several texture descriptors.

- Integrate toolboxes into user interface

PARTIALLY COMPLETED – The integration of toolboxes has been under way. This integration will be completed when the toolboxes are fully developed.

KEY RESEARCH ACCOMPLISHMENTS

- Developed optimized software for complete 64-channel beamformer (see Appendix 1)
- Synchronized the beamformer with the transducer and pulsing system.
- Design receive electronics and test for a single channel using a prototype array
- Designed and evaluated a complete multi-channel transmit board
- Fabricated and evaluated a prototype array (see Appendix 2)
- Developed interactive and dynamic multi-planar reformatting (> 30 frames/second)
- Demonstrated real-time scan conversion of simulated 4D ultrasound data
- Demonstrated rigid and nonrigid registration of ultrasound (see Appendix 3)
- Developed a multi-function optimization algorithm (see Appendix 4)

REPORTABLE OUTCOMES

Manuscripts, abstracts, presentations:

- C.R. Hazard and G.R. Lockwood, "Developing a High Speed Beamformer Using the TMS320C6201 Digital Signal Processor," *Proc. 2000 IEEE Ultrasonics Symposium*, San Juan, Puerto Rico, In Press.
- Vladimir Zagrodsky, Raj Shekhar, and J. Fredrick Cornhill, "Multifunction extension of simplex optimization method for mutual information based registration of ultrasound volumes," to be presented at the SPIE Medical Imaging Symposium in San Diego in February 2001.
- Raj Shekhar, Vladimir Zagrodsky, and J. Fredrick Cornhill, "Mutual information-based rigid and nonrigid registration of ultrasound volumes," submitted to *IEEE Transactions on Medical Imaging*.

Funding applied for:

- Development of quantitative real-time 3D stress echocardiography (PI - Raj Shekhar)
Funding Agency: The Whitaker Foundation
- Improving coronary artery disease diagnosis through registration of real-time 3D and SPECT images (PI - Raj Shekhar)
Funding Agency: American Heart Association – Ohio Valley Affiliate

CONCLUSIONS

We have developed a method for high speed ultrasound imaging. The method permits the acquisition of 40 to 80 images in the time normally required to collect a single image. The increased acquisition speed will be used to collect a 3D data set in real-time. The objective of this proposal is to build a prototype scanner and develop the software tools required to analyze and interpret the 3D data.

The beamformer is integral to allowing high speed ultrasound imaging and collection of 3D data sets in real-time. Using the design derived from our previous simulations and modeling, we have assembled the hardware to implement this beamformer. During this second year, we have fully developed and optimized the assembly language routines required for high speed beamforming. We have moved from a theoretical evaluation to a real world implementation. This required overcoming the additional challenges of real-time communication between a large network of processors, as well as synchronization of all the components in the system. While some testing remains, we are confident that the beamformer will be capable of real-time 3D data acquisition.

In order to allow our synthetic aperture approach to imaging, the transducer array must have an improved signal-to-noise ratio over conventional array technology. Tetrad Corp. has developed a multi-layered array for this purpose. A prototype of the array has been fabricated and used to design the front end electronics of the system. This multi-layer device provides increased sensitivity by improving the electrical matching between the transducer and the electronics. Since the multi-layer technology is being newly developed, there have been some challenges in perfecting the fabrication techniques for this array. Additionally, a rocking probe has been designed, and a prototype array tested in the fixture. A working prototype of the complete probe with a transducer array meeting our design criteria is expected in the beginning of the coming year.

The user interface and image analysis software development has focused on developing interactive visualization and image analysis algorithms. We have demonstrated real-time scan conversion of the native polar data for image preview and transducer placement. This software is being interfaced with the beamformer. The development of an interactive and dynamic MPR viewing mode for simultaneous viewing of multiple imaging planes of 4D data has been completed.

3D image segmentation and registration are two core image analysis technologies we have developed. These image analysis algorithms, together with the developed visualization tools, are fundamental to the creation of the cardiac and breast imaging toolboxes. Efforts are under way to acquire trauma images at a military hospital to complete the development of shrapnel detection toolbox.

Real-time 3D imaging combined with tele-medicine and a set of image analysis tools will enable ultrasound imaging for forward echelon combat casualty care. This technology will be equally effective in civilian emergency care.

APPENDICES

Appendix 1: Preprint of the publication - C.R. Hazard and G.R. Lockwood, "Developing a High Speed Beamformer Using the TMS320C6201 Digital Signal Processor," *Proc. 2000 IEEE Ultrasonics Symposium*, San Juan, Puerto Rico.

Appendix 2: Progress reports from Tetrad Corporation.

Appendix 3: Manuscript undergoing peer-review - Raj Shekhar, Vladimir Zagrodsky, and J. Fredrick Cornhill, "Mutual information-based rigid and nonrigid registration of ultrasound volumes," submitted to IEEE Transactions on Medical Imaging.

Appendix 4: Abstract of the paper to be presented at SPIE Medical Imaging Symposium in San Diego in February 2001 - Vladimir Zagrodsky, Raj Shekhar, and J. Fredrick Cornhill, "Multifunction extension of simplex optimization method for mutual information based registration of ultrasound volumes."

Appendix 1: Preprint of the publication - C.R. Hazard and G.R. Lockwood, "Developing a High Speed Beamformer Using the TMS320C6201 Digital Signal Processor," *Proc. 2000 IEEE Ultrasonics Symposium*, San Juan, Puerto Rico.

Developing a High Speed Beamformer Using the TMS320C6201 Digital Signal Processor

C. R. Hazard^{1,2} and G. R. Lockwood^{1,3}

¹Department of Biomedical Engineering, The Cleveland Clinic Foundation, Cleveland, OH.

²Biomedical Engineering Program, The Ohio State University, Columbus, OH.

³Department of Physics, Queens University, Kingston, Ontario, Canada.

Abstract--We are developing a synthetic aperture beamformer for a realtime 3D imaging system. The beamformer consists of a network of multi-processor digital signal processor (DSP) boards which provide the processing speed and flexibility required for realtime 3D beamforming. The system uses high speed memory buffers and a ribbon cable interface between DSP boards to allow channel to channel transfers to occur in parallel. Each channel has a 12 bit A/D, a C6201 DSP which calculates delayed signal values using linear interpolation, and a DSP for summing. Dividing the summing and delaying tasks reduces the input/output (I/O) burden. Restricting the A/D sampling frequency to an integer multiple of four times the center frequency of the ultrasound signal further reduces the I/O burden. An algorithm optimized for the C6201 calculates a single demodulated point in 25 ns. A multi-resolution look up table algorithm produces an 8 bit estimate of the log compressed envelope signal in less than 25 ns per point. Simulations of the radiation pattern formed using three transmit pulses show secondary lobes below -50 dB.

I. INTRODUCTION

We are developing a prototype realtime 3D imaging system based on synthetic aperture beamforming. This beamformer must calculate an image in the time that a conventional beamformer would calculate a few lines. Digital signal processors are an ideal platform for developing a beamformer. They provide the processing speed required for realtime beamforming, but are still easily adapted to different algorithms.

Previously, we introduced a beamforming architecture which accurately delays signals using a

linear interpolation algorithm, and sums values channel to channel using a pipelined network of DSPs [1]. Here we further investigate this architecture. The system generates a synthetic aperture image from three transmit events. We evaluate both the accuracy of the beamforming and the performance of the system as implemented with a fixed point DSP (TMS320C6201, Texas Instruments, Dallas, TX). An important limitation on the performance of a DSP based beamformer is the speed at which data can be moved between processors. In fact, for a high speed beamformer, the input/output (I/O) performance is as important as the CPU speed. We present algorithms directed at minimizing the I/O demands on the system. We also address the issue of demodulation and log compression using fixed point processors.

II. OVERVIEW OF PIPELINE ARCHITECTURE

The beamformer consists of a network of digital signal processors. Summing is done by arranging the DSPs in a pipeline so that the input to a processor is the sum of all the prior channels and the output is the sum including the contribution from that channel. Pipelining reduces the number of processors by eliminating the need for a hierarchical structure in which pairs of channels are summed by additional processors. However, pipelining introduces a data skew because the Nth processor cannot add data until all the processors up to and including the (N-1)th processor have calculated sums. This means that the Nth channel must either wait to calculate its values or store those values until the previous channels have calculated the sums. Waiting to calculate the values is constrained by the amount of memory available to store the digitized signals which would be needed for the calculations. Our system calculates the values and

stores the results until they are ready to be summed. The concept of skew is illustrated in figure 1.

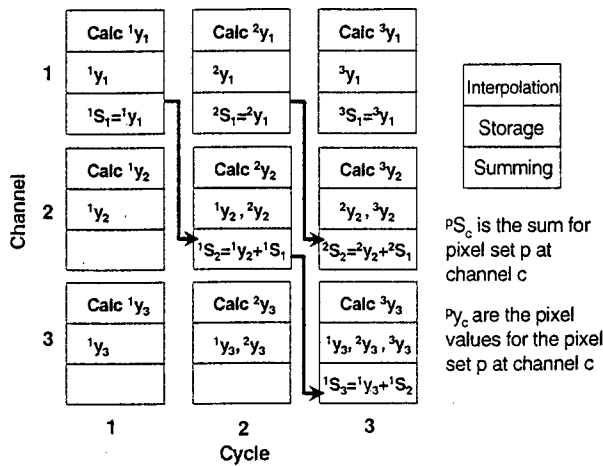


Figure 1: Data flow for pipelined system

The pipeline structure requires that data be transferred from one channel to the next. For our system, data is received at nearly 85 MB/second/channel. Interpolated outputs can be calculated at 160 MB/second/channel. These large data rates preclude the use of a standard bus architecture, since a bus would not allow the parallel transfer of data between channels. Instead we use 100 MHz, 32 bit first in first out memory buffers (FIFO) and a 160 MB/s ribbon cable interface which directly connects channels. This allows the simultaneous transfer of sums. The system uses DMA engines to transfer data. Each DMA transfer has a setup time and therefore data is most efficiently transferred in large blocks which minimize the setup penalty. The data skew introduced by pipelining and the amount of memory available limits the size of the blocks. Therefore skew leads to a trade off between memory and I/O efficiency.

III. HARDWARE IMPLEMENTATION

We have implemented a linear interpolation synthetic aperture beamformer using the 200 MHz TMS320C6201 DSP. The system is based on a multiprocessor board manufactured by Pentek, Inc. (Model 9137, Upper Saddle River, NJ). Each board provides four C6201 processors; two channels of 12 bit, 42 MHz A/D; and two Front Panel Data Ports (FPDP) in a single VME slot. Each board houses two channels

consisting of an A/D and two C6201s. One DSP is used to calculate delayed and apodized values using linear interpolation [1], while the second DSP sums these values with the values from the pipeline.

Signals from the transducer are digitized at 42 MHz. Two A/D samples are packaged in a 32 bit word and then written into a FIFO at 84 MB/s. The data can be read from the FIFO by the C6201 at 400 MB/s. A complete synthetic aperture image is formed using information collected from three separate transmit events. Data for the first two transmits of an image are stored in the internal data memory. The third transmit is temporarily stored in the FIFO itself. This allows collection of all three transmits in a time that is limited by the speed of sound and not processing speed. By collecting the three transmits quickly we reduce the system's sensitivity to motion [2].

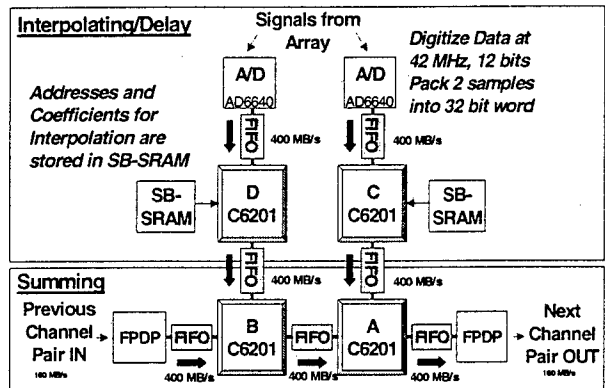


Figure 2: Hardware implementation of pipelined, linear interpolation beamformer

Figure 2 shows the architecture of a single multiprocessor board. The C6201, connected to each A/D, calculates apodized and delayed values using linear interpolation. In figure 2, the interpolating C6201s are labeled C and D. The interpolating C6201s also read addresses and coefficients from SB-SRAM. The interpolators calculate the delayed values and store them temporarily in the internal data memory. This temporary storage helps offset the data skew introduced by the pipeline architecture. The second C6201 provides additional skew buffers. After interpolation the values are transferred to a second processor which is responsible for summing the data. The second processor is part of the pipeline. Transferring data between processors is complicated

by the fact that the A/D FIFO, the SB-SRAM, and the FIFO connected to the second processor all share the same external memory interface. Consequently, these resources cannot be accessed simultaneously, despite the parallel nature of the CPU and DMA on the C6201.

The second processor for each channel adds the values for that channel to the sum of the values from the previous channels in the pipeline. On figure 2, the adders are labeled A and B. Each adder also has three connections to the external memory interface of the C6201: a FIFO input from the interpolating C6201, a second FIFO input, and a FIFO output. For processor B, in figure 2, the second FIFO input is from the 160 MB/s ribbon cable interface that connects multiple boards. The ribbon interface provides the sum of the values from the previous channels on other boards. Processor B calculates a new sum by adding the values from processor D, which were temporarily stored in skew buffers. This new sum is then sent to processor A using the output FIFO for processor B. For processor A, the second input FIFO is connected to processor B and contains the newly summed values. Processor A adds the values from processor C and sends the sum to the next board using the ribbon cable.

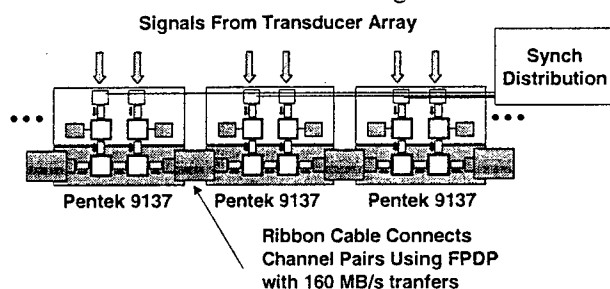


Figure 3: Multiple board connections

Figure 3 shows how the multiple boards are interfaced using the ribbon cable interfaces. The A/Ds from each channel are synchronized by distributing clock and start signals to each board. The high voltage pulser that excites the transducer array is also synchronized to the A/D's. The total system has 64 channels, consisting of 32 VME based Quad C6201 boards (two channels per board). An additional board is used to provide DC offset corrections, sum data over multiple transmits, perform envelope detection calculations, log compression, and interface with a PC for scan conversion.

IV. SOFTWARE IMPLEMENTATION

The software was written in hand optimized assembly language to ensure high performance beamforming. There are several design considerations for programming the C6201. The internal data memory of the C6201 is only 64 kB. The CPU has two sets of arithmetic/logic units which allow up to eight instructions to occur simultaneously (though many restrictions apply). The external memories are more efficiently accessed using the DMA and large blocks. The CPU pays a large overhead for each access to external memory. All external memory is accessed through the external memory interface, which precludes parallel access [3].

Proc A	Proc B	Proc C	Proc D
Read Delayed Values from C	Read Delayed Values from D	Simultaneously	Simultaneously
Read Partial Sums from B	Read Data from Previous Board	1) Read A/D 2) Read Coeffs from SBSRAM 3) Delay with linear interp.	1) Read A/D 2) Read Coeffs from SBSRAM 3) Delay with linear interp.
SUM Values	SUM Values		
Write Data to Next Board	Write Partial Sums to A	Write Delayed values to A	Write Delayed values to B

Figure 4: Main tasks for each processor

Figure 4 shows the main tasks for each processor. The algorithms are synchronized to allow realtime beamforming. This synchronization is achieved using matched algorithms and limited handshaking. The beamforming software is designed to minimize I/O, which is the limiting factor. The I/O necessary for the linear interpolation algorithm can be minimized by restricting the A/D sampling frequency to an integer multiple of the center frequency of the transducer array. This ensures that there are samples a quarter period apart. We estimate the signal envelope by calculating inphase (I) and quadrature (Q) components of the signal for each point in the image. The quadrature component is estimated by introducing an additional quarter wavelength delay, relative to the inphase component. The envelope is then estimated by $\sqrt{I^2 + Q^2}$. Restricting the sampling frequency ensures that the linear interpolation coefficients are the same for both I and Q. In addition only a single address needs to be stored, as the I and Q sample values will always be

separated by a fixed number of samples. Using this sampling scheme we can pack the delay address and the linear interpolation coefficients for I and Q into a single 32 bit word. This greatly reduces the I/O between the C6201 and SB-SRAM.

Packing the I and Q values into a single 32 bit word for transfer reduces I/O and allows the use of a special add instruction on the C6201, which sums the I and Q in a single cycle. However, this restricts the summing accuracy to 16 bits, which limits the dynamic range of the system.

The beamforming assembly code has been optimized by carefully arranging buffers to avoid memory resource conflicts, and when possible the calculations are done in parallel with DMA data transfers. The linear interpolation routine calculates packed IQ pairs in 25 ns. The summing algorithm sums IQ pairs in 8 ns. The system is capable of calculating 6.5 million IQ pairs per second.

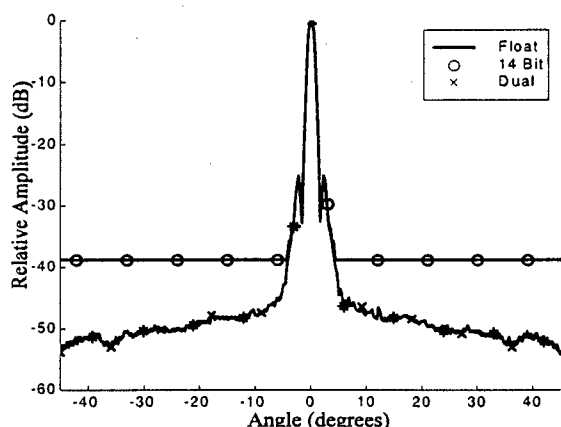


Figure 5: Radiation patterns using floating point, 14-bit LUT and dual LUT

A final DSP board at the end of the pipeline is responsible for envelope detection and log compression. For a fixed point processor, direct calculation of these complicated functions would be too slow for realtime beamforming. Instead the system must use a lookup table. We evaluated the performance of the lookup table by comparing the resulting radiation patterns to that produced by a floating point calculation. Figure 5, shows the radiation pattern of a point target at $f/4$ for the system using floating point math, using a 14 bit lookup table, and using a dual lookup table. A 21 bit lookup table ensures that the radiation pattern for the fixed point

system is less than 1 dB different than the floating point calculation. However, a 21 bit lookup table would require over 2 MB of memory. A 14 bit lookup table requires only 16 kB, but severely limits the dynamic range of the system. The secondary lobes for a 14 bit table are down by less than 40 dB. A good compromise, which meets both our accuracy and memory requirements, is a dual lookup table. The dual table uses a 14 bit lookup table for all but the lower 1000 values, for which a portion of the 21 bit table is used. This provides an 8 bit estimate of the log envelope in less than 25 ns per point. The memory required is less than 17 kB and the accuracy is less than 1 dB different than the floating point calculations (as seen in the figure). The secondary lobes are below -50 dB for the system.

V. SUMMARY AND CONCLUSIONS

We have presented the design for a beamformer capable of producing over 6.5 million points per second. The beamformer is limited by the I/O capacity of current state of the art DSP hardware. However, by restricting the sampling rate of the A/D to reduce coefficient storage, packing data for efficient transfer, and carefully managing memory resources this bottleneck can be reduced. The system is capable of calculating IQ pairs in 25 ns and summing IQ pairs in 8 ns. A dual resolution lookup table provides realtime envelope detection and log compression. This system can be used as a high speed beamformer for a synthetic aperture imaging system.

VI. REFERENCES

- [1] C. R. Hazard and G. R. Lockwood, "Theoretical assessment of a synthetic aperture beamformer for real-time 3D imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Contr.*, vol. 46(4), pp. 972-980, 1999.
- [2] C. R. Hazard and G. R. Lockwood, "Effects of Motion on a Synthetic Aperture Beamformer for Real-time 3D Ultrasound." *Proc. 1999 IEEE Ultrason. Symp.*, vol. 2, pp. 1221-1224, 1999.
- [3] "TMS320C6000 Technical Brief." TI Literature # SPRU197D, February 1999.

Appendix 2: Progress reports from Tetrad Corporation.

Project: Cleveland Clinic 3-D Mechanically Steered Phased Array Probe
Customer Grant #: DAMD17-99-1-9034
Report: Decision on single vs multilayer ceramic, delivery of phased array
Milestones: #4 and #6
Date: July 7, 2000
Written By: Mike Zipparo
Organization: Tetrad Corporation
Address: 357 Inverness Dr. South, Suite A
Englewood, CO 80112
Phone: 303-754-2309
FAX: 303-754-2329
E-mail: mzipparo@tetradcorp.com

Introduction

This report summarizes tests conducted on multilayer ceramic purchased from outside vendors and internally at Tetrad. It was found that multilayers purchased externally did not stand up well to dicing and nor maintain performance when placed into an array. Multilayers manufactured at Tetrad were found to perform much better, including in array form. Phased arrays were manufactured at Tetrad and tank tests were conducted to determine the ability of the transducers to withstand high voltage drive. For comparison purposes results for two multilayer phased arrays were analyzed along with a commercial single layer phased array.

The tank measurements were performed with the transducers being driven by a unipolar negative polarity pulse. The pulser delivered 278 V peak into an open circuit and 270 V peak into a 50 ohm load. The voltage delivered to the complex transducer impedance was either higher or lower, as noted for the appropriate test. Since the pulser voltage was of a fixed value, the only way to increase the power delivered to an element was to pulse at a higher rate, hence increasing only the temporal power but not the instantaneous power. Repetition rates between 500 Hz and 5000 Hz were employed to determine the effect of temporal power on the received amplitude. In all cases the transducer was driven with a 100 ohm resistor in parallel with the transducer.

The single layer phased array operated with a 4.0 MHz center frequency, while the multilayer arrays operated with center frequencies of 3.0 and 3.2 MHz. The thickness of the ceramic of the single layer array was 0.323 mm, while the thickness of the individual layers of the multilayer was 0.146 mm.

Executive Summary of Results

Two phased arrays were manufactured using Tetrad multilayer ceramic. Both arrays had 100% of their elements with full capacitance. Thus it is apparent that multilayers made at Tetrad are able to withstand the array fabrication process, including dicing into individual elements, better than commercially available multilayers. The arrays were tested while driven with a high voltage source (220 V) which resulted in a field close to the poling field. When driven in the direction of poling, pulses at that voltage and a repetition rate of 1.5 kHz, more than the maximum anticipated to be used in the final system, showed no degradation in signal level even after 2 hours. The signal level of the multilayer arrays was found to be significantly higher than a comparable array made with a single

layer of ceramic. Based on this outcome it is clear that proceeding with a multilayer for the final deliverable array is the best decision.

Experimental Summary

The single layer phased array showed signs of slight depoling at a high drive voltage, but the received signal amplitude was not lowered even at the highest rep rate of 5 kHz. The first multilayer phased array showed a higher received voltage when the drive signal was applied in the direction of poling. The received signal was not degraded even with a rep rate of 5 kHz for 23 minutes. The second multilayer array showed that driving in the direction opposite to the poling results in a lower received voltage due to depoling, and the rate of depoling is greater for higher rep rates. This depoling was found to be reversible. It was also found that it is possible to permanently damage the multilayer, even if it is driven in the direction of poling, if a high enough rep rate is used for an extended period of time. The damage is due to partial disbonding of the multilayer and is non-reversible. It was found, however, that if the rep rate is kept at 1.5 kHz or lower, a drive signal close to the poling field can be used for over 2 hours with no loss in received sensitivity.

The angular response of multilayer phased array #2 was found to be 69 degrees. Based on the preliminary measurements of the acoustic window in the wobbling mechanism, this should be sufficiently wide so that even with a small amount of beam narrowing due to the window the -6dB beamwidth should still meet the 60 degree criteria.

Previous results with multilayers

Multilayers were received from two vendors. The multilayers received from Vendor A were found to be unsuitable for ultrasound applications. The measured coupling was extremely weak and the elements did not withstand dicing. In addition, it appears that the internal electrodes had a very high resistance. Multilayers from Vendor B were found to be of better quality, but still not acceptable for an array with 100% functional elements. The coupling was reasonable, but not all of the elements withstood dicing. An array was made with one of these multilayers and it was found that the elements did not withstand high power drive. At this point it was decided that it was more feasible to develop multilayer ceramic technology at Tetrad so that all of the issues affecting performance could be addressed. Initial results showed that it was possible to manufacture multilayers in-house that yielded 100% functional array elements when diced into array elements. Work then proceeded with applying Tetrad triple layer ceramic multilayers into phased array designs so that acoustic performance could be evaluated.

Single layer phased array

The pulser voltage measured at the transducer terminals was 540 V. It is speculated that this is due to either the complex and large impedance (150 ohms at 3.7 MHz) of the single layer array element or due to the coax cable between the pulser and the transducer. The echo received from a flat plate with the transducer at the focal distance was 165 mV, for a loss of 70 dB. The signal decreased to 133 mV after 40 minutes with a 500 Hz repetition rate. Apparently this was due to the elevation angle shifting slightly, because by adjusting the angle the signal increased to 151 mV. After this the polarity on the drive signal was reversed and a signal of 147 mV was recorded. After 17 minutes the voltage reduced to 139 mV. At this time the rep rate was increased first to 1 kHz and then to 5 kHz, with neither affecting the signal amplitude appreciably.

Multilayer array #1

A scaled version of the 4.0 MHz single layer array was tested with the same setup, first with a 500 Hz rep rate. This phased array was built with a pitch scaled so that the elements were slightly wider, in terms of wavelengths, than the single layer array. The voltage measured at the transducer was 222 V and the received signal was 493 mV, for a loss of 53 dB. The loss numbers might not be directly comparable, since the effect of the 100 ohm resistor will be greater on the higher impedance transducer. Also, the elevation of the two arrays was the same; therefore the beam divergence and hence the loss of the lower frequency multilayer array will be greater. These effects should partially offset so that at least approximately the overall loss comparison should be valid.

After 32 minutes this signal reduced to 462 mV. With the polarity reversed the received echo was 511 mV, and after 9 minutes this reduced to 497 mV. Reversing the polarity apparently put the driving voltage in the same direction as the poling voltage, thereby reinforcing poling. At this point the rep rate was increased to 5 kHz. After 23 minutes the signal was 496 mV. Even a high rep rate did not result in a lower received voltage when the transducer was driven in the direction of poling.

Multilayer array #2

This phased array was built with a pitch scaled directly with frequency so that the element spacing was the same number of wavelengths as for the single layer array. The elevation was the same as the single layer array. The voltage at the transducer terminals was 220 V. The field across one layer of the multilayer was 15.1 kV/cm, close to the poling field of 20 kV/cm typically used for poling PZT ceramic. Three elements were tested. Element 1 was driven for 8 minutes at a 500 Hz rep rate. The received voltage went from 373 to 368 mV over that time period. At this point the rep rate was increased to 5 kHz, and after 34 minutes the signal was reduced to 96 mV. Apparently combination of the high field and rep rate caused the multilayer to depole significantly. This drive was in the direction opposite to the poling direction.

Element 2 was then driven at 500 Hz with a drive voltage at the transducer of 215 V. The measured receive voltage vs time is plotted in Figure 1. It can be seen that the received voltage decreased slightly over a 29 minute span with the rep rate at 500 Hz, continuing with a 1 kHz rep rate. The rate of signal loss increased with the 2.0 kHz and the 3.5 kHz rep rates. Finally with a rep rate of 5 kHz, the received voltage was reduced to just under half of the original signal. This element was then repoled at a field of 20 kV/cm in air at room temperature and the received voltage was rechecked. Figure 2 shows the received voltage to be decreasing with the input in the same polarity as the original over a 6 minute period. The polarity of the input signal was then reversed and the received signal increased to 439 mV. After increasing the rep rate to 1 kHz for 26 minutes, the signal increased to 458 mV. Apparently with the polarity of the input signal reversed, the polarization of the ceramic is reinforced. Because the excitation signal is a negative pulse, it would be necessary to reverse the poling direction so that the negative pulse drive would reinforce the poling. At a 1 kHz rep rate the amplitude of the received signal was maintained for over 25 minutes. It appears that the transducer is capable of operating at a field close to the poling field for an extended period of time at a repetition rate of 1 kHz.

Figure 3 shows the measured sequence for element 3, which was driven at 500 Hz in the original polarity and measured 388 mV on receive. The polarity on this element was reversed so that the

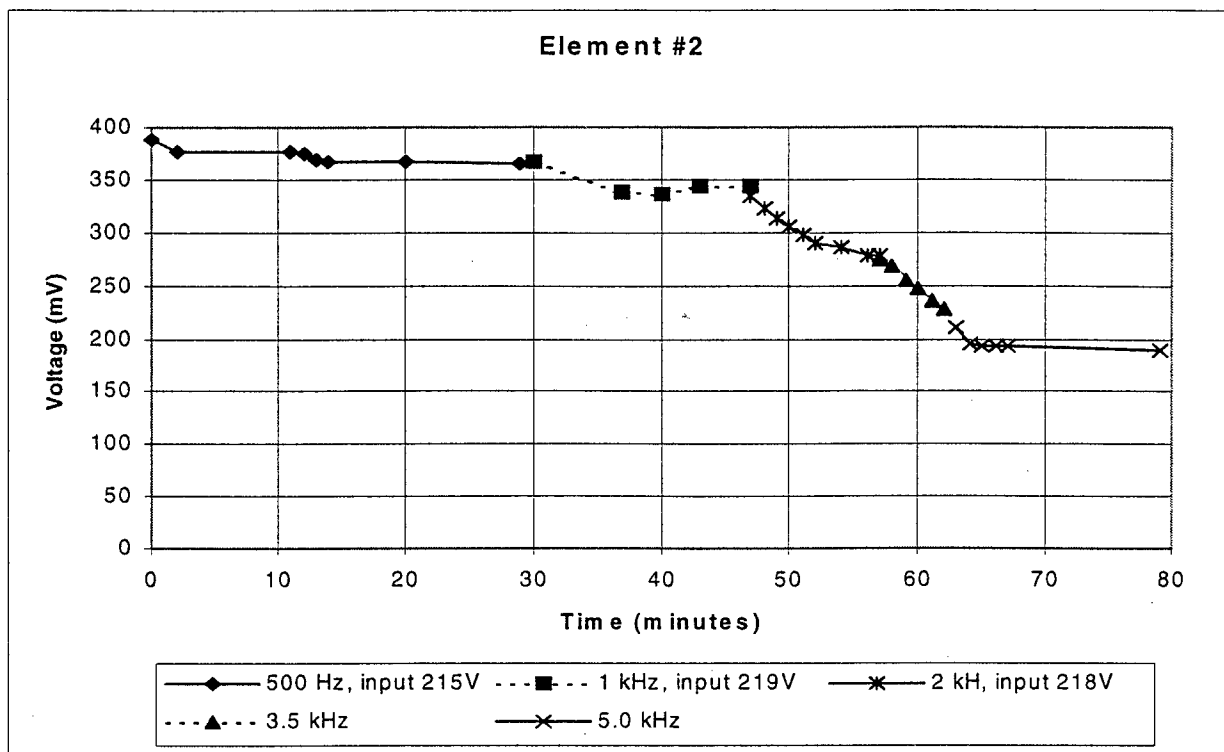


Figure 1) Received voltage vs time for 3-layer ceramic transducer #2, element #2 at various repetition rates with an input of 218 volts (15 kV/cm) oriented opposite to the poling direction.

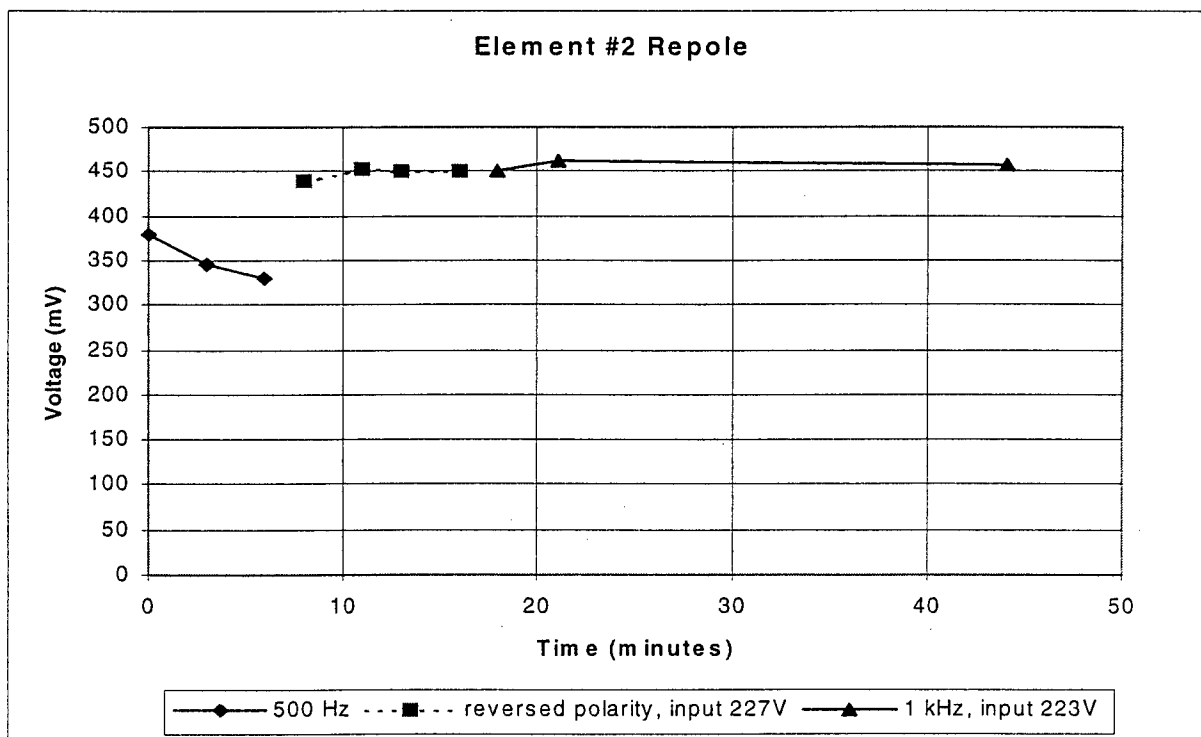


Figure 2) Received voltage from element #2 repoled and driven first opposite to the poling direction, then with it, showing that the initial loss of received signal was due to reversible depoling.

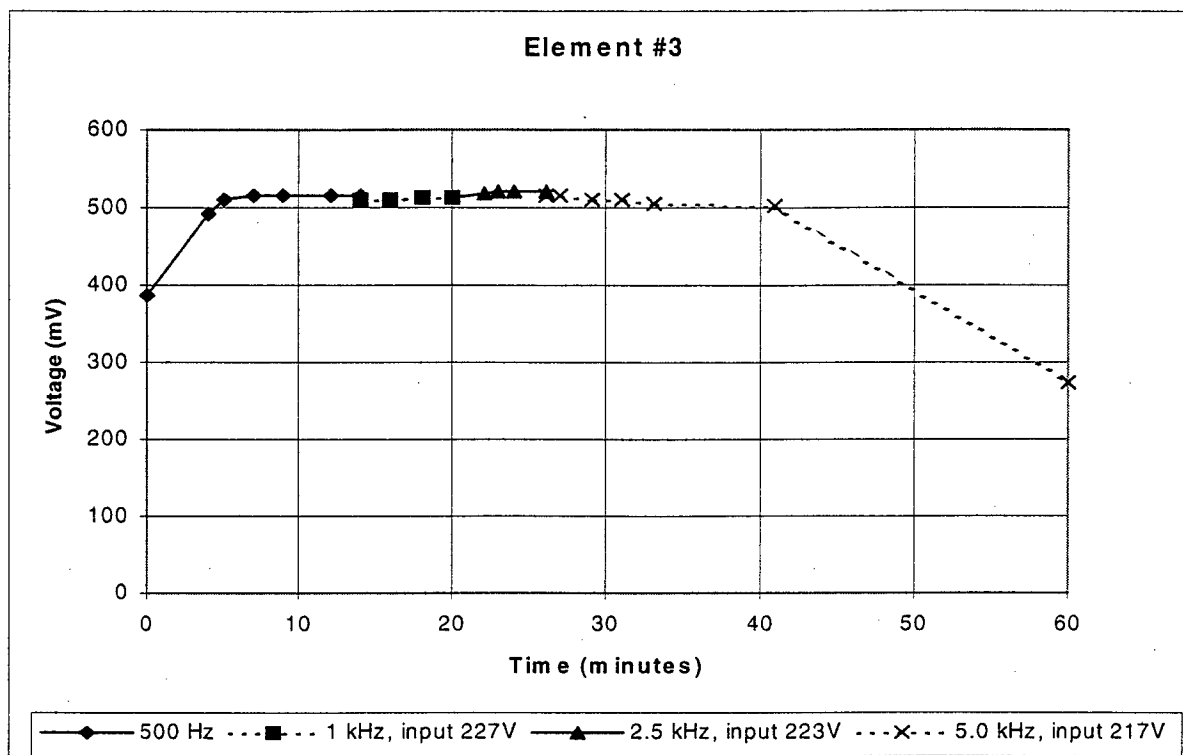


Figure 3) Received voltage from element #3, initially driven opposite to poling for 4 minutes, then driven with poling. Signal degradation after 40 minutes was due to non-reversible multilayer disbond, most likely from excessive heating in the piezoelectric ceramic.

drive signal reinforced the poling. It can be seen that even after 15 minutes at a rep rate of 5 kHz the signal was not degraded significantly. After leaving the transducer unattended for an additional 20 minutes, however, the amplitude was reduced significantly. Repoling of this element was unsuccessful. From a look at the electrical impedance for this element, it is evident that the multilayer had disbanded and this was the reason for the loss in signal. Apparently a high drive amplitude and a high rep rate for a long period of time generates enough heat for the multilayer to disband and the damage in this case is non-reversible.

An additional test was performed on the repoled element #2, first driving it at 500 Hz in the direction of poling for 8 minutes, then at 2 kHz for just over 2 hours. The results are plotted in Figure 4. After 1 hour the received voltage was almost the same, and after 2 hours it was reduced by 20 %. At this point the element was repoled again, but after subsequent testing the element did not receive signal did not improve. Analysis of the electrical impedance showed a small degree of disbonding in the multilayer, but not to the extent of element 3. Thus a 2 kHz rep rate seems to be enough to heat the multilayer to failure if driven for an extended period of time.

A fourth element was tested with the rep rate set at 1.5 kHz, and the results are shown in Figure 4. It is clear that even after driving for an extended period of time at a 1.5 kHz rep rate, the multilayer is still functioning effectively and without any noticeable degradation in performance.

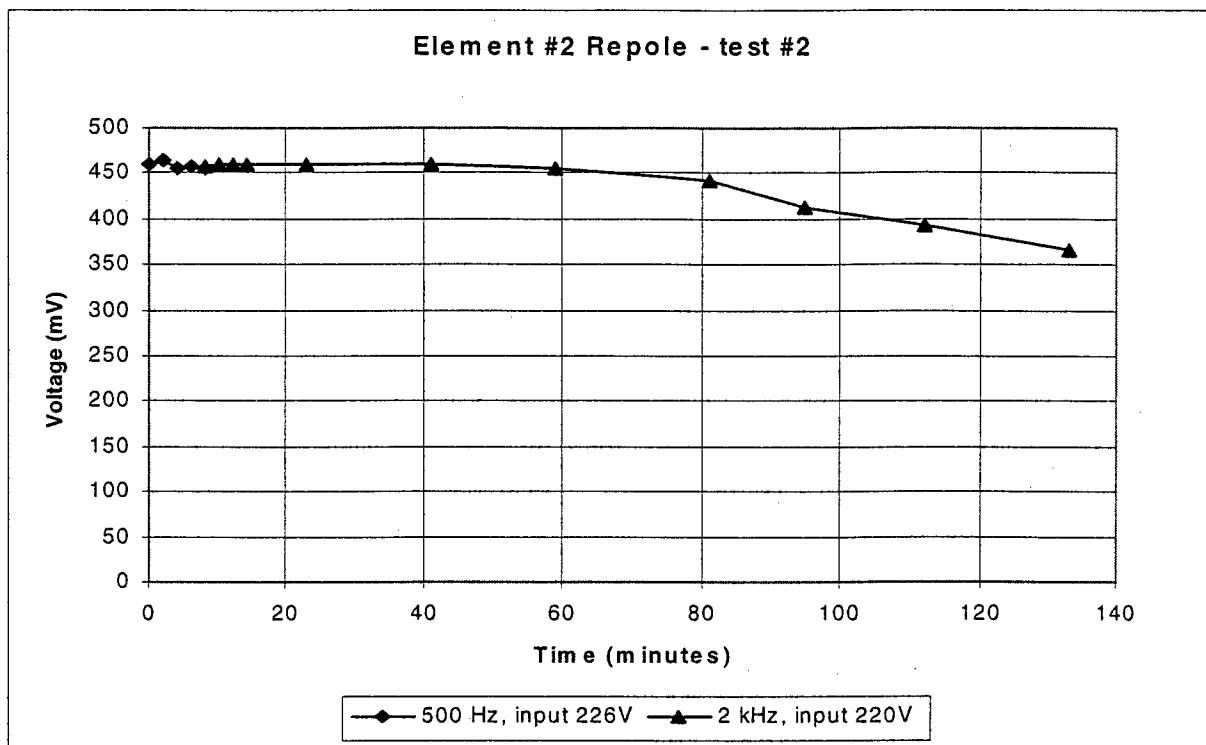


Figure 4) Received voltage from element #2 after repole, second test. Some signal degradation occurred after extended drive at a 2 kHz rep rate, due to non-reversible disbonding of the multilayer.

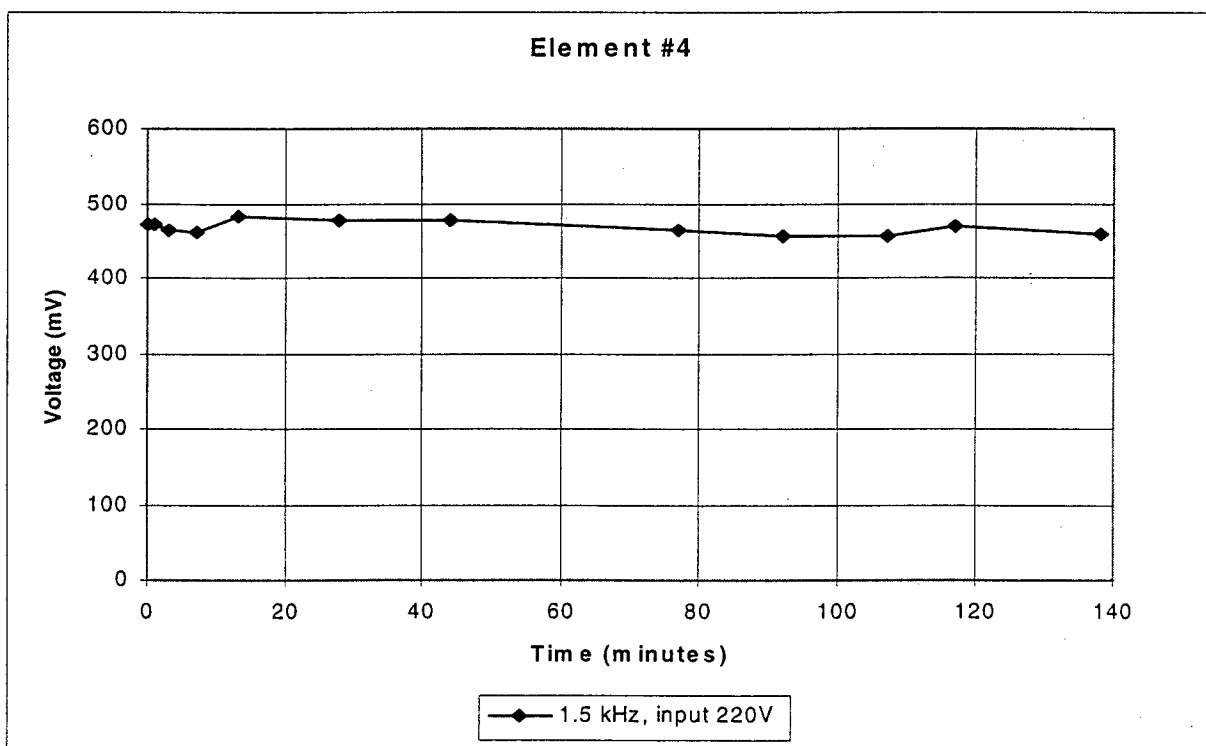


Figure 5) Received voltage from element #4. Some variation in received voltage was observed, but no significant degradation in received signal amplitude was observed after 2 hours of pulsing at a 1.5 kHz rep rate.

To: Geoff Lockwood
Project: Cleveland Clinic 3-D Mechanically Steered Phased Array Probe
Customer Grant #: DAMD17-99-1-9034
Reports: Present acoustic design and materials
Mechanical probe design with drive system specs, life time of flex
Feasibility, Phase I summary
Milestones: #5, #7 and #8
Date: August 21, 2000
Written By: Mike Zipparo
Organization: Tetrad Corporation
Address: 357 Inverness Dr. South, Suite A
Englewood, CO 80112
Phone: 303-754-2309
FAX: 303-754-2329
E-mail: mzipparo@tetradcorp.com

Introduction

This report serves as a final feasibility report on Phase I of the project, including the acoustic design and materials (array angular response, coupling fluid, and window material), mechanical probe design with drive system specs, and life time of flex. Other milestones under Phase I have already been reported on, including analysis of SNR, multilayers incorporated into phased arrays, phased array prototyping, available mechanical probe technologies, and other critical issues.

Summary of Phase I – Feasibility Study

The purpose of Phase I of this program was to initiate a design approach and conduct preliminary experiments to determine the feasibility of making a phased array in a configuration which will be capable of being rotated mechanically in one direction. Because of the low transmit power used in the imaging approach, the required sensitivity of the array needs to be greater than a conventional array. Previous reports have detailed the technical risks associated with the proposed approach for rotating the array. The decision to use a multilayer in the array has been made based on the successful construction of an array with 100% of the elements functional and verification that the elements remain functional at high transmit voltages.

This report summarizes the remainder of the feasibility study, including a theoretical analysis of the acoustic window and measurements on an initial prototype. Also included are results on the measured effect of the acoustic window on the beamwidth of the probe.

The effects of the acoustic window on the intensity of the transmitted beam and beamwidth are primarily due to attenuation in the coupling fluid and the window, rather than from reflection at any interface. The attenuation could be decreased by making the window thinner, though this would make the window more flexible and could hinder mechanical rotation in a realistic operating environment. The reverberation has been measured with the elements oriented normal to the window and found to be the expected amplitude. Measurements off axis were unsuccessful because of the low acceptance angle of the linear array in the first prototype wobbler. The affect of near field artifact due to reflection will probably have to be assessed

using the Cleveland Clinic system during early evaluations.

The longitudinal wave critical angle of the current window/fluid system is just beyond 45 degrees. However, due to the velocity of the fluid the worst case center ray of from the phased array (45 degree steering angle) is incident upon the window at only 41 degrees. Therefore all imaging can be done at less than the critical angle. There will be time delays and attenuation that vary with the incident angle as discussed below. These are an inevitable consequence of scanning through a fluid-plastic interface at varying angles unless properties can be matched exactly.

A complete array in the wobbler mechanism was tested for reliability. No interconnection failures occurred after continuous rotation for over 100 hours.

Acoustic design and materials

This section reports on the geometry and materials of the acoustic window, including measurements of velocity and attenuation for the actual window materials and their effect, including transmission and reflection on the beam as it is steered from zero through 45 degrees.

Acoustic window geometry and material properties

The phased array is attached to rotation mechanism (wobbler) contained inside a fluid-filled enclosure with a window material on the front, as shown in Figure 1. The axis of rotation is located at the center of the window curvature so that the array maintains a constant distance from the window as it is rotated. The window material is polystyrene with a thickness of 0.49 mm, and the coupling fluid is raw linseed oil. Both materials were tested to determine their acoustic properties, and the results are shown in Table 1.

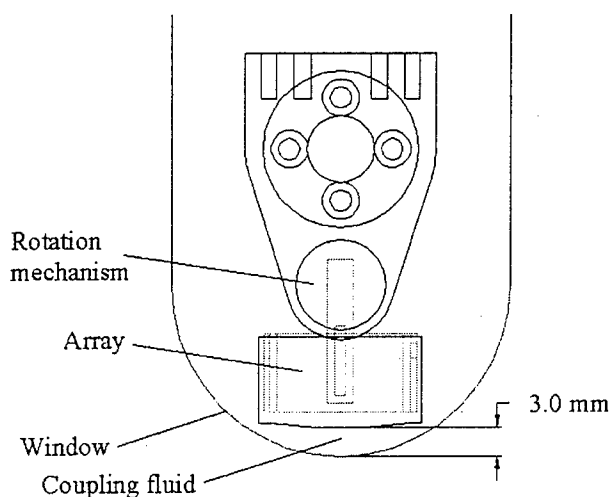


Figure 1 – Drawing of the phased array in the wobbler mechanism

Transmission loss through the window

A theoretical analysis of the transmission characteristics of this arrangement has been conducted, and the results are described below, including a comparison of measurements made on the physical probe. This analysis assumes that both media are liquid, an assumption that neglects the

conversion of waves from longitudinal to shear and vice versa. It also assumes that the waves are

Table 1 – Acoustic properties of window material and coupling fluid

		polystyrene	linseed oil
Velocity	(mm/us)	2.00	1.42
Density	(g/cc)	1.01	0.92
Impedance	(MRayl)	2.02	1.33
Attenuation	(dB/mm)	3.3 *	0.6 **
* measured at 3.5 MHz			
** measured at 5.2 MHz			

normal to the surface.

Using the above assumptions, the power loss from reflection at the interface between the coupling fluid and the window is 0.18 dB, and between the window and water is 0.09 dB. Thus the total round trip loss from reflections is 0.54 dB. Based on the measurements, the loss from attenuation is much greater, with a total round trip loss of 3.3 dB in the window and 3.6 dB in the fluid. The loss in the fluid was measured at 5.2 MHz, and should be appreciably lower at 3.2 MHz, though we could not measure the loss accurately at that frequency. This loss increases with increased steering angle due to the angle that the beam takes through the fluid and the window. The path length through the window at a 45 degree angle is approximately 2.6 times what it is for an unsteered beam (See Figure 2 below). This is from refraction at the fluid-window interface, and that will have a significant effect on the attenuation at large angles, since the loss in the window is almost six times as great as in the fluid. At 30 degrees, the path length in the window is only 30% longer than for zero degrees. There will also be an increase in path length in the fluid due to the geometry of steering the beam at an angle. Because of the lower attenuation in the fluid, this will be less significant than the attenuation in the window. Based on the total path length in the fluid and window, the net loss at zero, 30 and 45 degrees is 6.9, 8.5, and 13.6 dB. The loss calculated at 30 degrees agrees well with measurements made on an actual probe (see *Angular response* below).

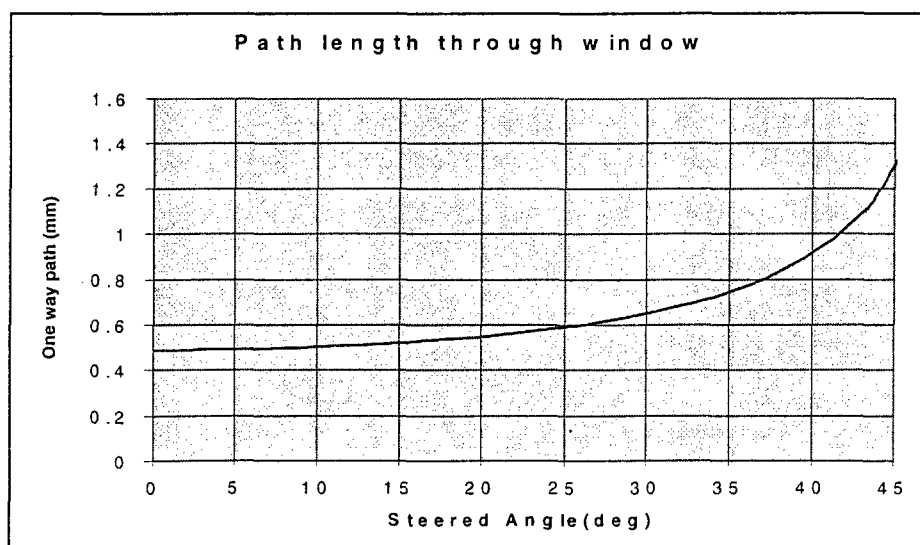


Figure 2 – One way path length from the transducer to the outside of the window as a function of steering

angle.

Refraction effects

Analysis of the refraction affects are slightly more complicated. Because there is a difference between the velocity of the fluid and that of the human body there will be a net refraction of the beam when it leaves or enters the probe. A plot of the incident angle in the probe versus the beam angle relative to the line normal to the center of the window is given in Figure 3. Note that the full 45 degrees of steering will be achieved when the incident angle between the fluid and the window is approximately 41 degrees. This should be transparent to the beam steering system which will still use delays appropriate to the human body. However, it is important in evaluation of the critical angle because the incident angle in the probe will be less than the steered angle in the body. It is also important because of the relative placement of line locations in the scan converter.

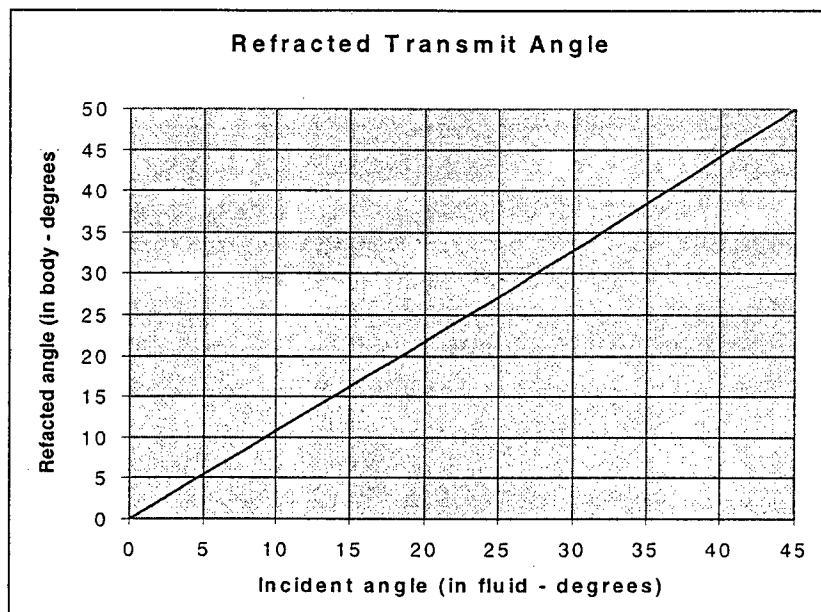


Figure 3 – Plot of refracted angle in the body vs. incident angle in the fluid. This plot shows that for an incident angle of 41 degrees on the fluid, the beam will be refracted to 45 degrees in the body.

Transmission and reflection through the acoustic window

A plot of the calculated transmission and reflection coefficients for the fluid-window interface (ignoring shear conversion) are given in Figure 4. Note that the current selection of acoustic materials produces a critical angle very close to 45 degrees. The transmission at 41 degrees, however, is still very good. Measurements made with the transducer indicate that the transmission coefficient at 45 degrees (in water) is still substantial.

The time delay caused by the window is a function of angle even without shear conversion. This is caused by the refraction of the longitudinal wave when it reaches the window. The combination of the difference of beam angle in the fluid and the path through the window produces displacements of the rays. The length of the one-way path through the window as a function of steering angle in the body is given as Figure 2. This increase in length causes substantially higher attenuation of the signal at large steering angles.

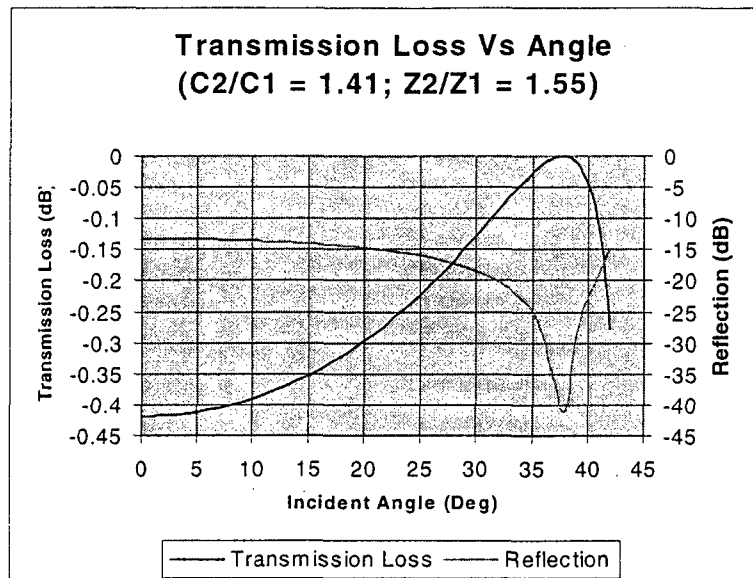


Figure 4 – Plot of the reflection and transmission coefficients for the fluid-water interface.

These affects could be reduced if the fluid and window could be perfectly matched but they cannot be. The critical angle could be increased either by selecting a fluid with a high sound velocity or a window with a lower sound velocity. It is difficult to raise the sound velocity of the fluid since it needs to be oil-based. Water-based fluids are not practical since the fluid needs to lubricate the probe movement. They are also more susceptible to bubble formation due to the variability of air solubility with temperature. Linseed oil has a relatively high velocity for an oil-based fluid.

Reverberation tests

Reverberation tests were performed on the wobbler mechanism containing a linear array. Multiple elements of the array were connected to a Panametrics pulser / receiver. First, two adjacent elements were tested with one of them as a transmitter and one as a receiver. Thus the incident wave is essentially normal to the window surface. The amplitude of the reflected signal using this configuration seemed reasonable.

As described above, the theoretical critical angle at which transmission through the window goes to zero occurs at just over 45 degrees. An attempt was made to determine the off-axis reflection characteristics by using elements at opposite ends of the array as transmitters and receivers. Unfortunately because of the limited angular response of the linear array, the amplitude of the signal was too low to measure at angles away from normal. Because of the close proximity of the array to the window material and the attenuation of the lens and coupling fluid, the reverberation should decay over a short period of time and would therefore cause artifacts only in the near field of an image. The net affect of this kind of reverberation is difficult to predict and will probably not be fully known until an image is created with the unit on the Cleveland Clinic system. It is thought that if reverberation is an issue with the phased array, then there is room in the wobbler mechanism for the placement of absorbers which could reduce their effect.

Angular Response

The angular response (FWHM) of the initial multilayer phased array was 69 degrees. Thus a small amount of beam narrowing due to the acoustic window can be tolerated and a 60 degree angular response maintained.

The effect of the window material on the angular response was measured by first testing the linear array outside the wobbler, and then testing it again inside the wobbler. The graphs are included at the end of the manual (see pages 21 and 22). From the graphs (which are normalized), it is evident that at 30 degrees, there is a reduction of just over 1 dB in the response with the window relative to the response without. The measured beamwidth of the phased array at -5 dB, which would correspond to -6dB trough the window, is roughly 62 degrees. If the angular response is required to be wider, it is possible to reduce the pitch slightly to increase the angular response, although this would decrease the sensitivity.

Acoustic window summary

There are a several materials that would produce a window with a lower sound velocity. All of these materials are all substantially more flexible and higher in attenuation than polystyrene. Flexibility presents a problem because thin membranes can easily be displaced by pressure on the face and could interfere with free transducer movement. Thick membranes present a problem because of the increased attenuation. The combination of a thicker membrane and higher material attenuation is likely to produce high differential attenuation at large steering angles relative to normal incidence even though the refraction path length through the window may be reduced.

It is our recommendation that this fluid and window system be used. If acoustic problems result we can run calculations to see if the problem observed might be improved by the use of an alternative window material. If performance can be substantially improved we can attempt a window replacement for the final phase of the project.

Mechanical Probe Design with Drive System Specs

An initial prototype of the wobbler mechanism has been constructed with a linear array and the acoustic window configuration described above. The complete manual for operating this unit is included as Appendix 1.

The unit was connected to a computer and the interface hardware was tested using the windows program "hyper terminal" as follows:

- 1) The probe control module was connected as shown in Figures 1 and 2 of the manual.
- 2) "Hyper terminal" was opened and a session started connecting directly to Com 1.
- 3) Under "properties" the settings were (see page 7 of the manual)
 - 9600 Baud
 - 8 data bits
 - no parity
 - 1 stop bit
 - no flow control

Note: It was necessary to go to "device manager" under "system" in "control panel" to install the driver for Com 1.

- 4) As capital letters, commands were typed according to the description on page 7.

Each of the scan presets shown in Table 4 were tested. The array elements were not connected during this experiment, as we did not have the appropriate system to image with this probe. The audio-frequency sound emanating from the probe was listened to as the setting was changed and a rotation initiated. As expected, it was apparent that the scan rate was increasing as the settings were changed from S1 to S6. The exact scan rate and sweep angle could not be measured without a more complex test setup. The developers of the mechanical system have thoroughly tested the accuracy and functionality of the sweep rates, and these can be checked through the AUX / IO signals if necessary.

Life Test

A complete wobbler mechanism with a linear array was tested to determine if it could withstand 100 hours of continuous cycling. The probe was tested with a ± 15 degree sweep angle at 20 Hz, after which time all of the elements were functional. In a previous test with a similar mechanism and flex circuit, the system ran for two months continuously with a ± 30 degree sweep angle (over two billion cycles). Given these results, it is unlikely that the lifetime of the probe will be an issue.

To: Geoff Lockwood
 Project: Cleveland Clinic 3-D Mechanically Steered Phased Array Probe
 Customer Grant #: DAMD17-99-1-9034
 Reports: Complete specification for 1st prototype
 Place orders for all components
 Milestones: #9 and #10
 Date: September 8, 2000
 Written By: Mike Zipparo
 Organization: Tetrad Corporation
 Address: 357 Inverness Dr. South, Suite A
 Englewood, CO 80112
 Phone: 303-754-2309
 FAX: 303-754-2329
 E-mail: mzipparo@tetradcorp.com

Milestone #9

A complete mechanical rotation mechanism has been built with a linear array contained inside. The results of this array, including response through the acoustic window and longevity of the flex circuit, have been described in the Phase I Feasibility Study. A multilayer phased array prototype has also been delivered under Phase I. The plan for Phase II is to deliver a second phased array with the same acoustic design and housed inside the wobbler mechanism which has already been built. The specifications for the wobbler mechanism are included as part of the operators manual which was included as part of the Phase I Feasibility Report. The specifications for the probe are as follows:

Probe type	1-D phased array in mechanical wobbler assembly
Mechanical specifications	See "3-D Probe System Specification" attachment
Window material	Polystyrene (same as initial prototype)
Coupling material	Linseed oil (same as initial prototype)
Center frequency	3.0 to 3.5 MHz
Elements	64
% working elements	100%
Sensitivity variation	Target -4 dB
Flatness	Target \pm 40 nsec
Connection longevity	20 Hz for > 50 hours
Piezoelectric material	PZT5-H, Tetrad bonded 3-layer ceramic stack
Active elevation	>13 mm
Pitch	0.27 mm (0.62 λ @ 3.25 MHz)
Angular response	
Target	69 degrees (same as initial prototype)
Min	60 degrees
Bandwidth	
Target	67% (same as initial prototype)
Min	60%
Power handling	Target – continuous unipolar drive at 220 V, > 1 kHz rep rate

Milestone #10

A process for building the initial multilayer arrays which will fit into the wobbler mechanism has been developed. The array will have the same acoustic geometry as the initial array prototype that has been delivered as part of Phase I. The only difference is that it will be built with a mechanism for mounting to the wobbler. All other array manufacturing processes will be identical to what has already been done. Tooling and parts which have been designed and ordered exclusively for this project include:

- 1) Custom flex circuit for interconnection to cable at array handle
- 2) Custom PC boards to match array pitch
- 3) Custom housing to mate with wobbler mechanism and to facilitate traditional array manufacturing process.
- 4) Custom tooling for holding housing during processing.

In addition, several multilayers for use in the final deliverable are in process. All of the other parts are either in-house already or can be fabricated in-house quickly for use when the custom parts are complete.

To: Geoff Lockwood
Project: Cleveland Clinic 3-D Mechanically Steered Phased Array Probe
Customer Grant #: DAMD17-99-1-9034
Reports: Build arrays for mechanical system – report
Milestone: #11
Amount: \$28k
Date: January 12th
Written By: Mike Zipparo
Organization: Tetrad Corporation
Address: 357 Inverness Dr. South, Suite A
Englewood, CO 80112
Phone: 303-754-2309
FAX: 303-754-2329
E-mail: mzipparo@tetradcorp.com

Milestone #11

A series of arrays have been fabricated according to the specifications listed under Milestone 10. One of the arrays has all of its elements working, although some at a lower level. This array has been tested extensively and will be put into the wobbler mechanism to test the acoustics of that system with a 3-layer phased array.

Figure 1 shows the capacitance measured for each element. There are 20 elements (out of 64 total) with an element capacitance which is approximately one-third of that of the remaining elements. The reason for this is that these elements were damaged during dicing and have lost a connection to one of their internal electrodes.

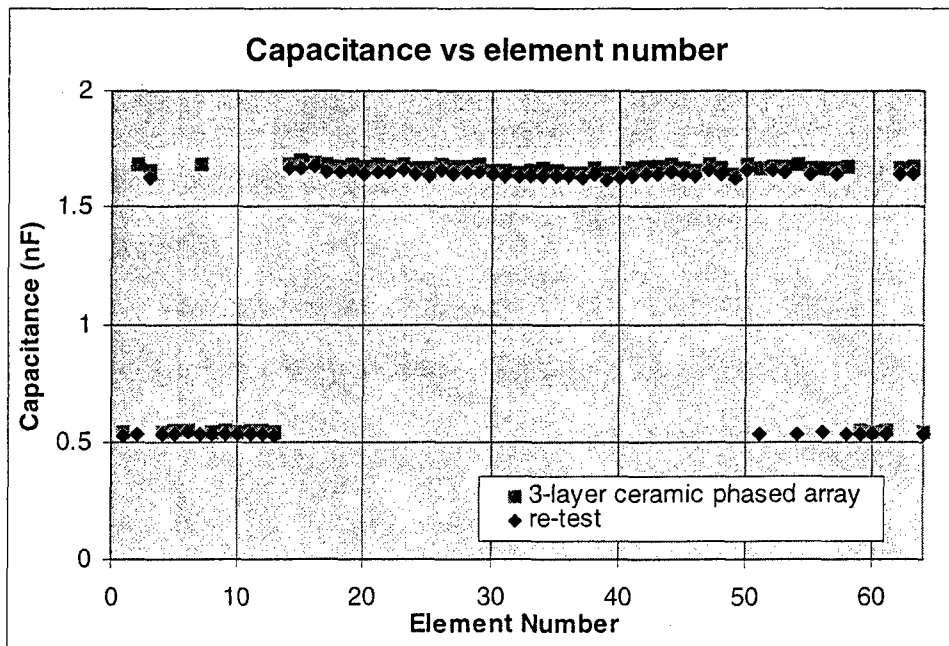


Figure 1) Capacitance vs. element. The diamond-shaped markers represent measurements just before sending the array to be placed into the wobbler mechanism.

The measured sensitivity is shown in Figure 2. It is apparent that most of the element which show low capacitance also show low sensitivity. This is because only one of the three layers is being driven electrically. The lowest elements are about 5 or 6 dB lower than the rest of the array. It should be noted that even with this level of decrease, the sensitivity of all of the elements is at least that which would be obtainable with a single layer of ceramic. The target sensitivity variation is -4 dB from the maximum, and this array has about -7 dB. The sensitivity uniformity will improve considerably when an array is built with all of the elements working with both internal electrodes connected. Technical difficulties are currently being addressed to achieve this objective.

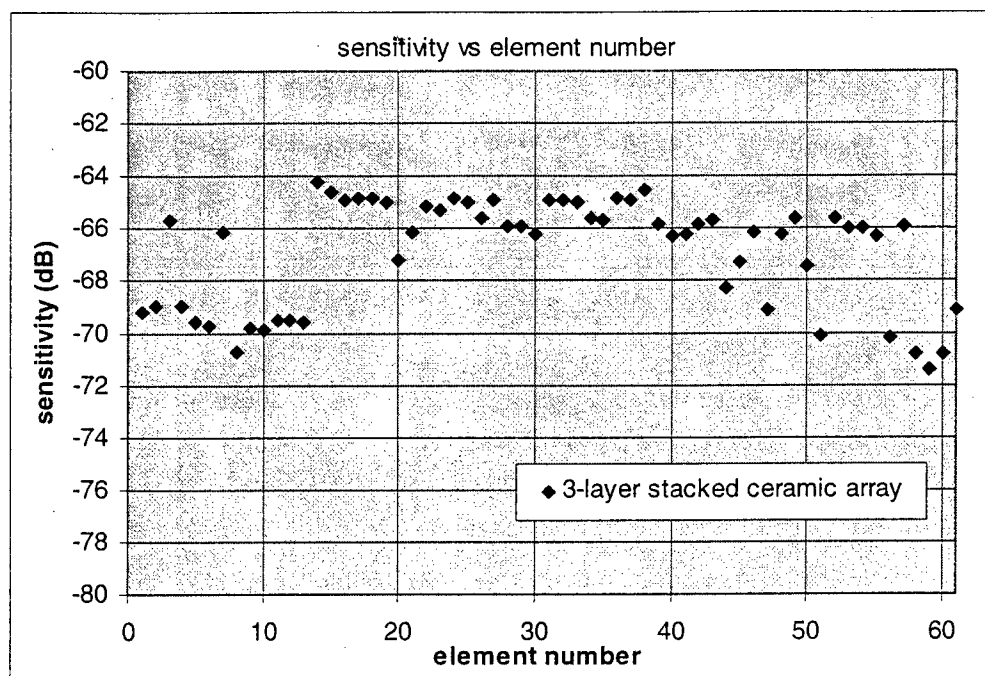


Figure 2) Measured sensitivity vs. element number. Most of the low sensitivity elements also measured low in capacitance.

The fractional bandwidth vs. element number is shown in Figure 3. The specification is a minimum of 60% bandwidth. The elements with low sensitivity also measured lower in bandwidth. This is because connection of only on-third of the multilayer results in a disrupted acoustic response. Even with the poorer elements, the mean bandwidth for the entire array was over 66%. Some element were over 80% bandwidth. Meeting the bandwidth spec should not be a problem when all of the elements are working properly.

The measured arrival time is shown in Figure 4. The absolute variation is 181 ns. However, a large portion of this is due to measurement error and misalignment of the probe during testing. Even if a probe has an arrival time variation of 40 ns, the test setup at Tetrad is not capable of measuring it.

The response vs. angle measured for element 24 is shown in Figure 5. The -6 dB level spans -34 to $+36$ degrees, for a total -6 dB beamwidth of 70 degrees. This is well above the specification

of 60 degrees.

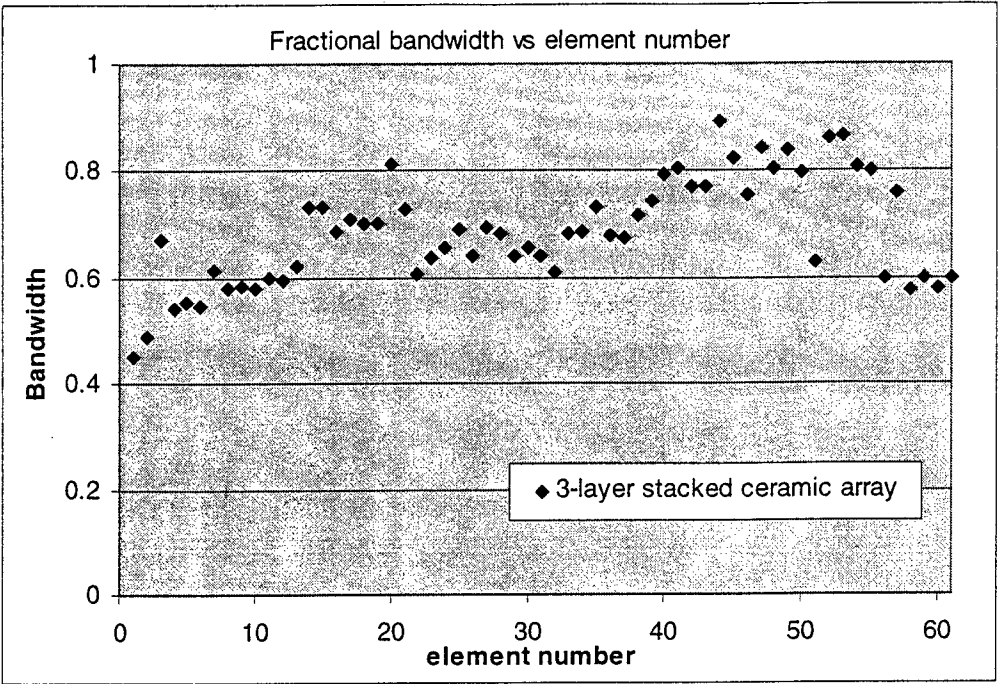


Figure 3) Measured fractional bandwidth vs. element number.

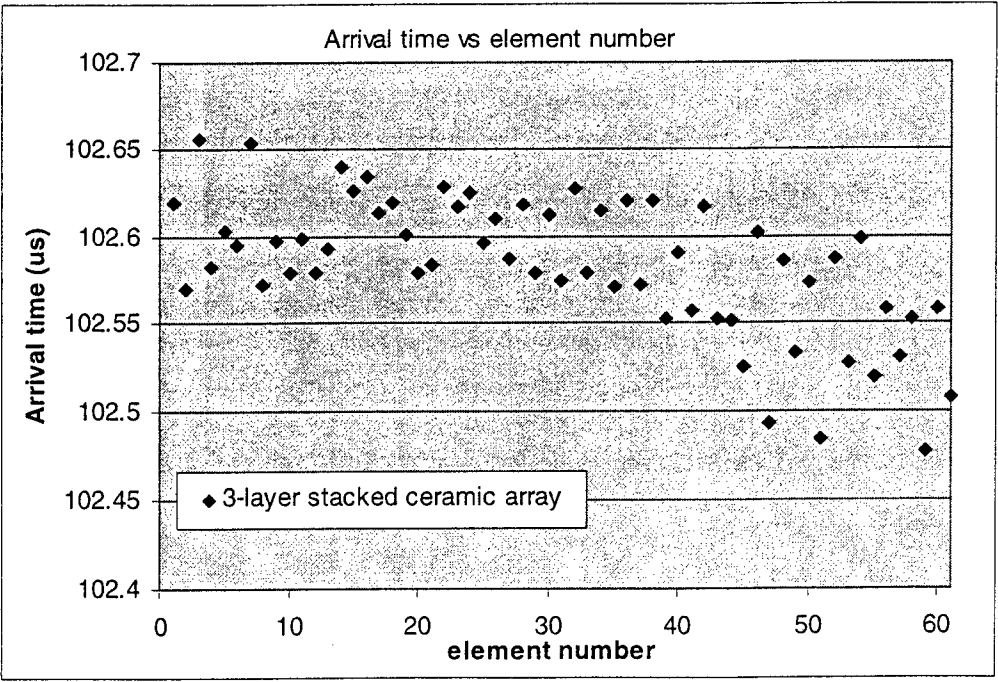


Figure 4) Measured arrival time vs. element number.

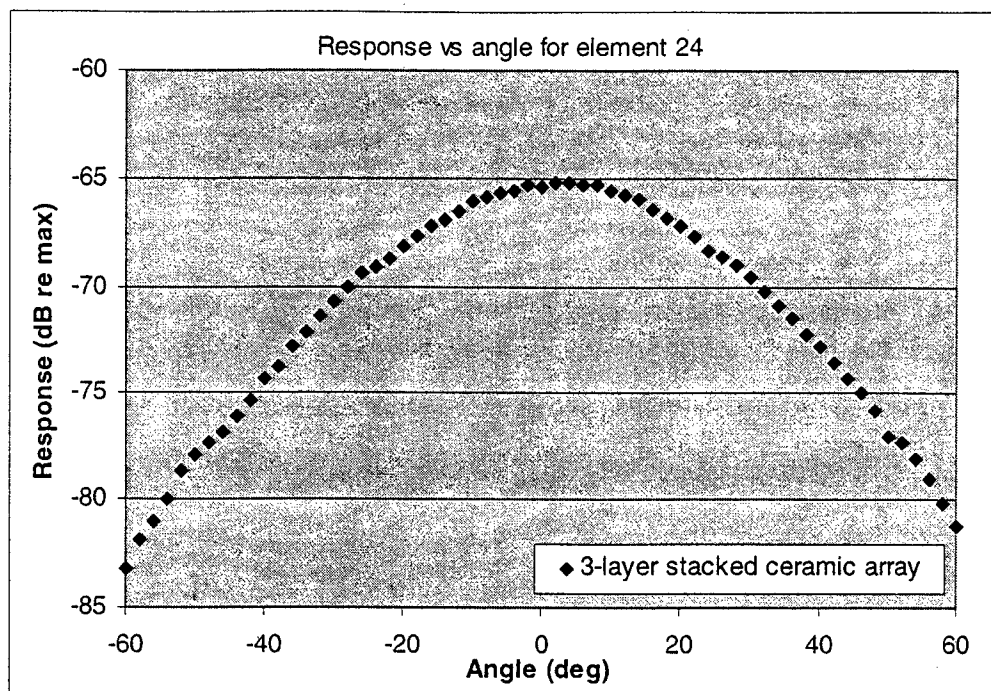


Figure 5) Measured angular response for element 24, showing a -6 dB beamwidth of 70 degrees.

Summary and Future Plans

An array has been built to fit into the mechanical rotation assembly. All of the elements are functional, although a significant number are operating at a reduced level because of damage done during the assembly. The probe as-is meets the requirements set for bandwidth and angular response, and the arrival time cannot be measured well enough to meet the 40 ns spec. Even with the bad elements, the sensitivity variation is still close to the spec of -4 dB. This will improve substantially when an array is made with completely functional multilayers.

This array will be placed into the existing wobbler assembly so that the entire acoustics system can be looked at with a phased arrays. Based on theoretical calculations made with the properties of the window system and on measurements made with a linear array in the housing, the current system is not expected to significantly reduce the -6 dB beamwidth. Following acoustics and life cycle tests of this array, a specification for the final array will be developed. This will conclude Phase II of our program.

The fabrication process is being examined closely to determine how to make an array with all of the layers working for all of the elements. The final Phase III probe is still expected to be a triple-layer ceramic with all of its elements fully functional. The drive system is already complete and will be delivered along with the Phase II array. Necessary improvements or modifications to the array or mechanical system will be made as necessary, based on the conclusions drawn from our tests and system incorporation and imaging done at The Cleveland Clinic Foundation.

Appendix 3: Manuscript undergoing peer-review - Raj Shekhar, Vladimir Zagrodsky, and J. Fredrick Cornhill, "Mutual information-based rigid and nonrigid registration of ultrasound volumes." Submitted to IEEE Transactions on Medical Imaging.

Mutual Information-Based Rigid and Nonrigid Registration of Ultrasound Volumes

Raj Shekhar, Vladimir Zagrodsky, and J. Fredrick Cornhill

Department of Biomedical Engineering, Lerner Research Institute
The Cleveland Clinic Foundation, Cleveland, Ohio 44195, USA

Corresponding Author:

Raj Shekhar, Ph.D.
Department of Biomedical Engineering (ND20)
Lerner Research Institute
The Cleveland Clinic Foundation
9500 Euclid Avenue
Cleveland, Ohio 44195

Phone: 216/445-3246
Fax: 216/444-9198
Email: shekhar@bme.ri.ccf.org

This project was supported by a Department of Defense research grant (DAMD17-99-1-9034, Principal Investigator – J. Fredrick Cornhill).

ABSTRACT

We investigated the registration of ultrasound volumes based on the mutual information measure, a technique originally developed for multimodality registration of brain images. A prerequisite for successful registration is a smooth, quasi-complex mutual information surface with an unambiguous maximum. We discuss the necessary preprocessing to create such a surface for ultrasound volumes. Abdominal and thoracic organs imaged with ultrasound typically move relative to the exterior of the body and are deformable. Consequently, four specific instances of image registration involving progressively generalized transformations were studied: rigid-body, rigid-body + uniform scaling, rigid-body + nonuniform scaling, and affine. Registration was applied to both clinically acquired and synthetically created images. The accuracy was comparable to the voxel dimension for the first two transformation modes. The accuracy degraded and the capture range shrank as the transformation grew more complex. Our principal objective was to demonstrate the feasibility of mutual information-based (i.e., segmentation-free) registration of ultrasound images of deformable organs. When real-time 3D ultrasound image acquisition becomes routine, this method should work well for to a variety of applications examining serial anatomic and physiologic changes. Developers of these clinical applications would match the deformation model of their problem to one of the four models presented here.

Key Words: image registration, mutual information, nonrigid image registration, three-dimensional ultrasound

I. INTRODUCTION

Registration of monomodality medical images is an important first step in successful visualization and quantification of temporal changes in anatomy and physiology. Since the bodily organs are fundamentally three-dimensional (3D), a comprehensive picture of serial change is expected to emerge from the registration of 3D images or volumes. A 3D approach is not only the most natural approach to medical image registration, it is also a worthwhile image processing endeavor.

Image registration has been an area of active research [1], and the state-of-the-art brain image registration solves many difficult clinical tasks [2-5]. However, there is a relative shortage of image registration work outside the brain anatomy, and consequently, a dearth of literature on registration techniques involving ultrasound imaging, a modality not suitable for the brain. Ultrasound, however, is ideal for imaging abdominal and thoracic organs, especially the heart. Even for such organs, registration of ultrasound images has not been reported in the literature much. This may be due to the relatively poor image quality of ultrasound and the nonrigid nature of organs typically imaged with it, to which registration techniques developed for the brain simply do not extend.

The lack of investigation in ultrasound image registration may also be due to the primarily two-dimensional (2D) nature of ultrasound. Whereas magnetic resonance imaging (MRI), computed tomography (CT) and various nuclear medicine image modalities have historically produced 3D images, ultrasound has not. There have been solutions suggested to reconstruct 3D by carefully registering 2D images acquired from a conventional ultrasound scanner. A 3D volume is reconstructed by translating, rotating or rocking the transducer head uniformly with the aid of a purpose-built mechanical device such that the spatial relationship between the acquired 2D images is known [6]. Alternatively, a 3D volume could be created by freehand manipulation of the transducer head whose orientation is recorded continuously with a wireless 3D localizer [7]. Regardless of transducer localization mechanism, such solutions are too slow to image a dynamic organ such as the heart. Electrocardiogram (ECG) and respiratory gating have been employed to reconstruct 3D images of the heart over multiple cardiac cycles, but such data may still have distortions due to cardiac arrhythmias and irregular breathing. Even when imaging a static organ, it is difficult to avoid distortion due to patient motion and inherent inaccuracies in 3D localization mechanisms.

Real-time 3D ultrasound acquisition, the most recent advance in ultrasound imaging, addresses both speed and distortion problems inherent in 3D reconstruction solutions. A commercial real-time 3D ultrasound system (Volumetrics, Inc., Durham, North Carolina) is available now, but at very few hospitals only. Where it is available, its use is limited to research investigations. Images are acquired through a 2D phased array of crystals, capable of directing an ultrasound beam anywhere within a 60 x 60 degree pyramid of space. Through the use of 16:1 parallel processing, it is possible to acquire 64 scan lines in 64 scan planes in less than 30 milliseconds, leading to effective volumetric imaging rates of greater than 25 Hz [8]. Although this scanner has the necessary speed for 3D image acquisition, it provides a lateral image resolution poorer than that of the current clinical images (64 vs. greater than 200 scan lines). Development is under way at our institution to produce a real-time 3D ultrasound scanner that works on the principle of

synthetic aperture beamforming [9, 10]. This scanner will not compromise lateral image resolution for speed and produce volumetric images with the current clinical lateral resolution. It is not difficult to envision real-time 3D acquisition as the future of ultrasound imaging, and most hospitals, especially their cardiology departments, adopting real-time 3D ultrasound in the near future.

There are several approaches to image registration, not all of which are applicable to our problem domain – registration of volumetric ultrasound images of the heart and other characteristically nonrigid organs. The anatomy and unique motion of the heart place special constraints on the registration approach. As commonly applied to brain image registration, frame-based techniques [11] or techniques that rely on the placement of external markers on the patient's body [12] assume a rigid underlying anatomy and a fixed spatial relationship of this anatomy with respect to the outside markers. The heart is not only nonrigid, it can move significantly within the chest cavity, rendering such approaches inappropriate. These prospective techniques, in general, have little clinical acceptability because they involve time-consuming acquisition protocols. Retrospective image registration, on the other hand, is nonobtrusive to the existing clinical practice and perhaps the only alternative in the case of abdominal and thoracic organs. Retrospective registration approaches utilize internal anatomic point, contour and surface landmarks, or voxel similarity [1]. Techniques based on internal landmarks, although generalizable to nonrigid transformation, have limitations because they involve some form of image segmentation. Not many point landmarks in ultrasound images of abdominal and thoracic organs can be reliably identified and used for registration. On the contrary, the requirement for the number of point landmarks is even higher to solve for nonrigid transformation. Contour and/or surface-based approaches rely on accurate segmentation of one or many anatomical structures in the images to be registered. Segmentation of ultrasound images is a difficult problem that usually requires manual intervention for optimal robustness and accuracy. If manual steps are involved, the accuracy of segmentation becomes user-dependent and is always suspect. Segmentation-based registration is limited by the accuracy, reliability and speed of segmentation. A voxel similarity-based approach provides the best framework for ultrasound volume registration. No segmentation of points, contours or surfaces is required, thereby removing any limits on the accuracy and speed. There is also no theoretical limitation on the nature of transformation (rigid or nonrigid) involved. A voxel similarity-based technique has the potential for full automation – another reason for its selection in the present investigation.

We here report results of registration of ultrasound volumes using mutual information measure [13] of voxel similarity. In particular, we report the accuracy, capture range and execution time for four different modes of possible transformation between two image representations of an anatomy. Although the results are presented for cardiac images, the approach is generalizable to ultrasound images of most other anatomical sites.

II. RELATED WORK

Voxel similarity-based techniques of image registration, especially those involving ultrasound images, form the background for our work. The flexibility of using voxel similarity for image registration has been recognized in the literature [1]. The superiority of a volume-based over a surface-based approach for multimodality registration of brain images has been shown [14].

Many studies [15, 16] have compared various measures of voxel similarity and concluded that mutual information is the most accurate and robust measure for 3D image registration. While these comparative studies have been performed on non-ultrasound data (brain MRI and single photon emission computed tomography (SPECT) in [15], liver MRI in [16]), we believe based on these studies and a preliminary study by our group [17] that the same holds for ultrasound. These studies provide sufficient confidence that mutual information is a reasonably good measure of voxel similarity. Instead of repeating experiments comparing voxel similarity measures, we have focused on preprocessing, recovery of rigid and nonrigid transformations, and, in general, effectiveness of the voxel similarity approach applied to ultrasound volume registration.

One of the first applications of voxel similarity for registering ultrasound volumes was by Rohling *et al.* [18]. The objective was to register up to six different freehand swept, volumetric ultrasound images of the gall bladder from slightly different viewing directions so that they could be spatially compounded (averaged) to create a 3D image free of acquisition distortions, artifacts and speckles. The specific voxel similarity measure used was the correlation coefficient of gradient images. The authors report the effectiveness of the voxel similarity approach and the adequacy of rigid-body registration to eliminate any artifacts due to organ movement between image acquisitions.

A study with a similar focus as ours is that by Meyer *et al.* [19], who used the mutual information measure successfully to register 3D ultrasound images of the breast. The objective was to register a pair of color flow and/or power Doppler images to create a difference image for serial monitoring of patients in response to chemotherapy or radiation therapy. If imaged from multiple directions, superimposition followed by registration allowed filling in of the flow information missing in a single view. Starting with an approximate registration based on user-selected point landmarks, the investigators refined the registration using voxel similarity. Rigid-body, full affine and elastic transformations were compared. The authors concluded that the affine transformation modeled the deformation of the breast between scans with clinically acceptable accuracy.

III. REGISTRATION METHODS

In this section, we briefly describe the general theory of mutual information-based 3D image registration and explain the ultrasound specific processing steps we have introduced.

A. Mutual Information-Based Registration

The algorithm assumes the existence of two data sets: one (primary) is kept stationary, and the other (secondary) is transformed iteratively until the most optimal alignment between the data sets, corresponding to the maximum of mutual information function, is achieved. An optimization method searches for the mutual information maximum in the domain of transformation parameters. The mutual information $I(A,B)$ between two data sets A and B is a function of the individual probability density functions $p(a)$ and $p(b)$ and the joint probability density function $p(a,b)$ of voxel intensities in the overlapping zone of A and B .

$$I(A, B) = \sum_a \sum_b p(a, b) \log \left(\frac{p(a, b)}{p(a)p(b)} \right) \quad (1)$$

Physically, mutual information conveys the amount of information that A contains about B , or vice versa [13].

In our formulation of the problem, the goal of image registration is to obtain a 4x4 transformation matrix T_0 , in homogeneous coordinates, such that the mutual information measure, $I(A, TB)$, between the primary set (A) and the transformed secondary set (TB) is maximized at $T = T_0$. T_0 refers to rigid-body transformation if it incorporates rotation and translation only. The transformation is affine if it includes scaling and shearing as well.

B. Modes of Transformation

A generalized affine transformation (T) is the cumulative effect of a series of scaling (S), shearing (H), rotation (R) and translation (D). Although individual transformations could be combined in many ways, we have restricted the order to the following.

$$T = D \times R \times H \times S \quad (2)$$

The expanded affine transformation matrix appears as below

$$T = \begin{bmatrix} r_{xx} & r_{xy} & r_{xz} & d_x \\ r_{yx} & r_{yy} & r_{yz} & d_y \\ r_{zx} & r_{zy} & r_{zz} & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad (3)$$

where $\{d_x, d_y, d_z\}$ is the translation vector, and the nine elements of the upper-left 3x3 submatrix encompass the combined effect of three rotations $\{\phi_x, \phi_y, \phi_z\}$, three scalings $\{s_x, s_y, s_z\}$ and three shearings $\{\theta_{xy}, \theta_{yz}, \theta_{zx}\}$ (refer to the Appendix for further formulation).

In the present work, we have investigated four different global transformation modes with progressively increasing complexity. The simplest mode corresponds to rigid-body transformation, whereas the most complex is full affine transformation. The two intermediate modes are limited forms of affine transformation. The first limited affine mode corresponds to rigid-body transformation plus uniform global scaling, whereas the second relaxes the uniformity condition to nonuniform scaling along the three principal axes. Shearing is not allowed in either mode. Totally elastic deformation, in which each data sample of the secondary set has a unique transformation, is conceivable, but such a registration may warp a pathological region of tissue loss or growth to align with a region of healthy tissue, defeating the purpose of serial follow-up by subtraction in medical applications.

The transformation mode determines the dimensionality of the search space for optimization. Six parameters, three translations $\{d_x, d_y, d_z\}$ and three rotations $\{\phi_x, \phi_y, \phi_z\}$, are searched in the rigid-body (RB) transformation mode. The uniform scaling (RB + US) mode searches for seven parameters – a global scaling parameter $\{s_x: s_y = s_z = s_x\}$ in addition to six transformation parameters of the RB mode. Three distinct scaling parameters $\{s_x, s_y, s_z\}$, one per principal axis, make the number of parameters searched in the nonuniform scaling (RB + NS) mode equal to nine. The last case, affine transformation (AT) mode, involves 15 geometric parameters that include three translations, three rotations, three scalings and six shearings. The effect of the 15

geometric parameters is, however, expressed by only 12 algebraic parameters in the 4x4 homogenous matrix formulation (see Eq. 3). In the Appendix, we have shown that only three shearing parameters are unique; the effect of the other three is a combination of the other geometric parameters. Therefore, in the AT mode, three shearing parameters $\{\theta_{xy}, \theta_{yz}, \theta_{zx}\}$ were searched in addition to nine parameters of RB + NS mode without any loss of generality.

C. Capture Range

Image registration using the mutual information measure searches for the maximum of mutual information function in the domain of transformation parameters. The desired solution, in general, is a strong local maximum but not the absolute maximum over the entire search space. This idea can be explained by an extreme example, in which, the two volumes overlap at only a few sample points. If the intensity histograms happened to be identical for the two volumes in the region of overlap, mutual information would be extremely high. However, such an orientation is clearly not the desired solution because the two volumes virtually do not overlap. An implicit assumption in our method, or any other voxel similarity-based method, is that the starting misalignment between the two data sets is in the neighborhood of the desired solution. We use the term *capture range* to convey the notion of a space around the perfect solution such that launching registration anywhere in this space guarantees convergence on the desired solution.

D. Speckle Reduction

A requirement for successful registration is that the mutual information function (or any other measure of voxel similarity) at least within the capture range must be quasi-convex with as few local maxima as possible. While this requirement is often met in mono- or multimodality registration of any combination of typical MRI, CT and nuclear medicine images, the speckle noise in ultrasound images makes this a difficult requirement to meet. Figure 1 shows surface plots of mutual information as a function of misalignment to illustrate this point. The mutual information surface for a pair of MRI and SPECT images in panel (a) is quite smooth, whereas the same surface for a monomodality 3D ultrasound image pair in panel (b) has ripples. Ripples or local maxima confound the search for the global maximum. Removal of speckles is key to suppressing the undesired local maxima in the mutual information function and making optimization robust and reliable. We accomplished this objective in two ways.

1) *Median filtering*: Speckles were reduced in a preprocessing step by a 3 x 3 x 3 median filtering kernel for clinical 3D ultrasound images. The kernel was 3 x 3 in size for semi-synthetic (to be explained later) data sets that were created from 2D image sequences of the heart. Median filtering suppressed speckles and in turn significantly smoothed the resulting mutual information function, as is apparent in panel (c) of Figure 1.

2) *Intensity quantization*: Sample points of ultrasound data are typically a byte or 8 bits long. Using fewer than 8 upper voxel intensity bits attenuates both signal and noise, but the signal-to-noise ratio (SNR) seems to improve first before decreasing, as the number of bits employed goes from 8 to 1. The higher the SNR, the smoother is the mutual information function. In the preliminary investigation [17], we showed that using either 5 or 6 upper bits allowed the most reliable convergence. Panels (c) and (d) of Figure 1 have the mutual information surface plots for 8 and 6 bits of intensity quantization, respectively, following

median filtering. The surface is smoother and steeper at 6 bits of intensity quantization. All results reported in this study are averages of repeating registration with both 5 and 6 bits of intensity quantization. We are investigating multifunction optimization that takes advantage of multiple levels of intensity quantization [20].

E. Optimization Algorithm

There are several approaches to optimization. Gradient-based approaches, although fast, are sensitive to errors in gradient calculation and the presence of local maxima. Gradient-based search algorithms were ruled out in the present study because of the limitation of the preprocessing to eliminate completely any speckle noise-induced artifacts from the mutual information function. The simulated annealing approach, despite its robustness, was also discarded because of the excessive time it took to converge. Based on our preliminary experimentation, we selected the simplex method of Nelder and Mead [21] as a compromise between robustness and convergence time.

Choosing the size of the initial simplex in a multidimensional parameter space is an important step in using simplex optimization. Each axis of the multidimensional parameter space corresponds to a geometric transformation parameter. For RB transformation mode, the space is 6-dimensional with d_x , d_y , d_z , ϕ_x , ϕ_y and ϕ_z as parameters. Prior to determining the size of the initial simplex, a normalization is desired such that a unit step along any parameter axis results in approximately the same physical displacement of the data in the spatial domain. The relationship between translation and physical displacement of data is direct – a unit translation moves all voxels of a 3D image by the same fixed amount. The same is, however, not true for rotation, where the displacement of a voxel is dependent on its distance from the axis of rotation. The physical displacement associated with rotation was estimated by the excursion of the farthest vertex from the axis of rotation passing through the center of the data. The displacements associated with scaling and shearing were similarly the excursions of the farthest vertex. In the present work, a physical displacement on the order of voxel dimension was chosen as the normalizing distance. The required translation (in millimeters), rotation (in degrees), scaling (unitless) or shearing (in degrees) to produce this physical displacement was treated as one unit of that parameter in the transformation parameter domain. For the data sets used in the study, a unit parameter distance corresponded to 1.25 mm of translation, 1 degree of rotation, 2% of scale and 2 degrees of shear.

In principle, the size of the initial simplex should be greater than the unit dimension along each axis so that it does not get stuck in a local maximum. The ripples in the mutual information function are spaced approximately at unit dimension. The initial size should also not be greater than the capture range, otherwise the convergence may not occur. An initial size roughly 2-5 units along each axis was found satisfactory.

When using optimization, one must also decide termination conditions. We employed a two-part condition; meeting both parts stopped the optimization. The first part checked for the size of the simplex. If it became smaller than a unit hypercube in the parameter space, the condition was considered met. The second part looked for the range of mutual information values at simplex vertices. This condition was considered met if and when the range became less than 0.001. There was also a failsafe condition, simply the number of iterations, which was empirically kept at

1000 to prevent optimization from executing indefinitely. The failsafe condition was rarely encountered, as the search would end due to the physically meaningful first condition.

IV. EXPERIMENTAL METHODS

A. *Description of Data*

There were two sources of 3D ultrasound data. The first was the real-time 3D ultrasound scanner manufactured by Volumetrics, Inc. This scanner produced a sequence of volumes, shaped approximately like a truncated pyramid, with 60 degrees azimuth and elevation angular spans, at a frame rate of 25 Hz. The scan depth was 140 mm for the data sets used in the present study. Volumetrics data sets were acquired natively on a spherical grid with a higher spatial sampling rate along the radial direction (referred to as the z axis henceforth). The number of samples along the z axis was less than 512; the actual number depended on the scan depth and the length of the null space coinciding with the near field of the ultrasonic beam. There were 64 samples along azimuth and elevation angles, the lateral sampling rate consequently varied with depth. To preserve the relatively higher spatial resolution along the z-axis, the data sets were scan-converted to a 128 x 128 x 512 rectilinear grid before further processing. The rectilinear grid coincided with the bounding box of the spherical grid. Five data sets (three from one patient and two from another) showing the left ventricle at different phases of the cardiac cycle were used. These clinical data sets were acquired in the Department of Cardiology at our institution.

Another type of 3D data set, which we call semi-synthetic, was created from time-series 2D images of the heart to simulate the type of volumetric data the real-time 3D ultrasound scanner that our group is developing will produce. All 2D images spanning a full cardiac cycle were spatially arranged to mimic the orientation resulting from the internal, user-transparent rocking motion of the transducer array, effectively converting time into a synthetic spatial axis. A cross-section of the resulting volumetric data parallel to the synthetic axis resembled a warped version of the familiar M-mode image. Although the semi-synthetic data did not represent realistic anatomy, such data closely approximated the higher lateral spatial resolution data that our scanner will provide. Since the success and accuracy of registration, to a great extent, is anticipated to depend on image resolution rather than the exact nature of the structure of the underlying anatomy, we believed semi-synthetic data would provide a closer estimate of the performance expected when our future images are registered. Semi-synthetic data sets originated from two image sequences that extended over three cardiac cycles each and showed the four-chamber view of the heart. The images were acquired in two patients as part of routine clinical diagnosis in the department of cardiology at our institution. The scan depth was 140 mm and the pixel size in the scan-converted 2D images 0.38 x 0.38 mm. Since the heart rate was stable, the number of 2D images per cardiac cycle was the same.

B. *Data Preparation*

The original 128 x 128 x 512 resolution of Volumetrics data sets overwhelmed the computer because 3D registration is computationally intensive. The sets were therefore subsampled by a factor of two to create 64 x 64 x 256 resolution data with a voxel size of 2.2 x 2.2 x 0.55 mm. A further subsampling along the z axis to create 64 x 64 x 64 resolution data sets was also performed for one of the experiments. In both cases, median filtering for speckle removal was performed. The overall data dimension was 140 x 140 x 140 mm.

To create semi-synthetic 3D data sets, a central square portion in each of the original clinical images was chosen; the surrounding annotation and ECG waveform were cropped out. Following 3×3 medial filtering, 32 consecutive cropped images (originating from one cardiac cycle) were stacked at an angular spacing of 1.8 degrees such that the apex of the underlying ultrasonic acquisition sectors met at a unique point. The resulting space of ultrasonic data was resampled to a $128 \times 64 \times 128$ rectilinear grid with a voxel size of $1.0 \times 1.9 \times 1.0$ mm, hence overall dimensions of $128 \times 122 \times 128$ mm. The poorer resolution corresponded to the synthetic spatial axis, i.e., y axis.

An important preprocessing step was to create a 3D mask to match the shape of the acquired ultrasound volume within the rectangular volume obtained following scan-conversion. The volume of overlap needed to compute mutual information was always the volume of overlap of the two masks. There was one unique mask for the Volumetrics data sets and another for the semi-synthetic data.

C. *Experiments and Validation*

We tagged, retrospectively, the frames of Volumetrics data sets with the cardiac phase and registered identical phase frames from two different scans of the same patient. The registration was found visually satisfactory in all cases; however, a quantitative validation could not be performed because the ground truth was not known. The experiments proved the feasibility of mutual information-based registration of ultrasound volumes. To determine the accuracy of registration of 3D ultrasound data sets for the four transformation modes and two types of 3D ultrasound data, we took a self-validation approach. Deforming the secondary volume in an otherwise registered pair of volumes simulated the starting misalignment. The goal of registration was then to overcome the user-introduced deformation by applying an exactly opposite transformation to the secondary. Comparing the known misalignment with the solution of registration allowed us to determine the expected accuracy. For each of the five Volumetrics data sets, two adjacent end-diastolic frames (separated in time by 40 ms) were chosen as the primary and the secondary sets. The proximity (in time) of the two frames and the end-diastolic phase, when the heart is momentarily stationary, allowed us to assume matching shape and size of the cardiac anatomy in the two frames. It further allowed us to assume the perfect anatomic alignment with partially, if not fully, uncorrelated noise. For each of the two semi-synthetic data sets, a 3D image pair was produced from two adjacent cardiac cycles. Since the heart rate was regular and the acquisition was fast, it was assumed that the deformation pattern of the heart was similar in the two cycles. Consequently, the two volumes were assumed to be registered initially without correlated noise.

We picked starting misalignment in the transformation parameter domain that allowed us to experiment with several different combinations of transformation parameters. We conceptually constructed a hypercube in the parameter domain centered at the origin with a side length of 12 parameter units. The number of vertices of the hypercube depended on the dimensionality of the space, which in turn depended on the transformation mode. It was 2^6 , 2^7 , 2^9 and 2^{12} , respectively, for RB, RB + US, RB + NS and AT modes. We picked 20 vertices of the hypercube at random for each transformation mode as 20 different starting misalignments. Each misalignment was then decomposed into its basic constituents (geometric transformation parameters) and applied to

the secondary before initiating registration. To summarize, each image pair for each transformation mode was registered 20 times starting from 20 different misalignments.

Upon registration, the known and the searched geometric transformation parameters were compared directly to assess accuracy. As discussed before, the actual physical displacement of each voxel of a 3D data set upon a complicated transformation (involving more than translation) is not identical. The physical displacement is usually greatest at the farthest corners or the vertices of the data volume. An alternative approach to estimate the accuracy of registration with just one metric was to consider the average error in the registration of the eight vertices of the data volume. Since the size of Volumetrics and semi-synthetic volumes differed, we averaged the errors at the vertices of a hypothetical cube of 100 mm side length centered with the data for a fair comparison of accuracy between the two kinds of data. This metric was termed *average distance error* and was measured in millimeters.

A subsequent goal of our experiments was to estimate the capture range for each transformation mode. The general idea was to plot average distance error against starting misalignment and look for a transition point beyond which average distance error increased dramatically. As before, we worked in the parameter domain. Conceptually, nine random rays were drawn originating from the origin. Starting misalignments were points along the ray between 4 and 15 units at unit interval. The capture range estimated in parameter units was converted to physical units.

V. RESULTS

A. *Effect of Median Filtering*

We discussed earlier the effect of median filtering on smoothening of mutual information function to facilitate convergence. In this section, we first present results demonstrating the effect of median filtering on the accuracy of registration for Volumetrics data sets. The average distance error, although increased with increasing complexity of deformation, was consistently smaller when the data sets were median filtered prior to registration (see Figure 2). Although not shown, the error in the estimation of each geometric transformation parameter was less with median filtering. All subsequent results include median filtering performed as part of preprocessing.

B. *Registration of Real-Time 3D Ultrasound Data*

The registration of one of the five Volumetrics data pair is shown in Figure 3. For each transformation mode corresponding to a row, two orthogonal cross-sections (x-y and y-z planes) of the fused volume data are presented. Each cross-section is presented twice, showing the relative orientation of the primary and the secondary before and after registration. The primary has been depicted with shades of green using the green channel of the RGB color triplet, whereas shades of magenta (red and blue channels) depict the secondary. Shades of gray result upon registration when comparable intensities of green and magenta are fused. A visual approach to evaluate the success of registration, therefore, is to look for a higher occurrence of gray. Qualitatively, a greater matching of anatomical structures is apparent following registration in all transformation modes.

Results of registration on all five Volumetrics data sets are shown in Table 1. The average distance error, averaged for all five data sets, is presented in the second column. This error increases with the complexity of transformation. The root mean squared (rms) deviation (or estimation error) of each parameter from the expected solution is reported in subsequent columns. The bottom number in each row is the deviation from the zero transformation, the obvious expected solution. Since there is bound to be some, although very small, mismatch in the two starting frames we used, the zero transformation may not be the ideal solution. As an alternative, we considered the median of all solutions (to be called median solution) as the expected solution and measured rms deviation of each geometric parameter from it. The numbers in the top of each row designate deviation from the median solution. Not surprisingly, the estimated parameters were closer to the median solution than the zero transformation solution. A general trend was increasing error in the estimation of each transformation parameter with increasing complexity of transformation. d_x (translation along x axis) is estimated with 0.41 mm rms error in RB mode, whereas the same error increases to 3.4 mm in the AT case. This observation is likely due to the ambiguity arising from a higher number of parameters with more complex transformation modes. A single basic transformation or a combination of them may approximate another single basic transformation. As an example, scaling along z axis may correct for translational error along the same axis in some parts of the volume. If the scaling is not allowed at all, as would be the case in RB mode, translation parameter will be more accurately determined. The anisotropic image resolution also has an effect in estimating transformation parameters. In general, any parameter that moved the data parallel to the axis of lower image resolution entailed more error.

One of our goals was to estimate the capture range for each transformation mode. A desirable way to express capture range is to define a range of values for relevant geometric parameters. However, geometric transformation parameters interact in a complex manner. A large translation could be recovered when present by itself, but the optimization may fail if the same translation is present along with rotation and scaling. An alternate way to determine capture range is in terms of parameter units encompassing the effect of all applicable parameters. The parameter unit can be converted to physical displacement for a more intuitive measure of capture range. In the four panels of Figure 4, the accuracy of registration (average distance error) is plotted against misalignment in parameters units. The boundary of capture range was estimated to be located at the point of inflexion where the slope of the fitted line changed. According to the preceding definition, the capture range was estimated to be 11, 10, 8 and 6 units, respectively, for RB, RB + US, RB + NS and AT transformation modes. These numbers translate to 14, 12.5, 10 and 7.5 mm in physical units, respectively. As expected, the capture range shrinks with increasing complexity of the transformation mode.

C. *Registration of Semi-synthetic 3D Ultrasound Data*

The images showing the result of registration on one of the semi-synthetic data sets are presented in Figure 5 in exactly the same layout as that of Figure 3. The success of registration is visually apparent by comparing pre- and post-registration images, especially the fused images following registration. As before, the accuracy is higher for simpler transformation modes. One can see that the algorithm did not achieve good registration upon convergence in the AT mode for the data set shown.

Table 2 presents a numeric summary of results for the semi-synthetic data sets. As before, the error in registration and estimation of a transformation parameter increased by increasing the complexity of the transformation mode. A difference is in the relative magnitudes of error between three forms of translation, rotation and so on. Due to lower sampling rate and the lack of closed features to constrain registration along the synthetic (y) axis, estimation of parameters that moved one volume with respect to another along the y axis was less accurate. As a result, the error in d_y was more than the error in d_x and d_z , and the error in s_y more than that in s_x and s_z . Since ϕ_y , the rotation about the y axis, affects x and z coordinates which were more finely sampled, it was more accurately estimated than ϕ_x , and ϕ_z that affected the y coordinate as well. By the same token, θ_{xz} was more accurately determined in comparison to the other two forms of shear.

Semi-synthetic data, in comparison to Volumetrics data, do not show dramatic improvement in average distance error. In fact, the error worsens in all cases except the RB + US case. The expected improvement was not achieved because the parameters associated with the synthetic axis could not be determined with high accuracy. The parameters not associated with the movement of data along the synthetic axis, d_x , d_z , ϕ_y , s_x , s_z and θ_{xz} , showed comparatively better accuracy. The average distance error significantly improved if the starting misalignment and the ensuing optimization left out the parameters associated with movement of data along the synthetic axis (see Table 3). This confirms a direct relationship between image resolution and the accuracy of registration.

D. Multiresolution Registration and Execution Times

In an attempt to test multiresolution strategy of image registration for Volumetrics data, the registration was attempted at several levels of resolution. The registration at the original data resolution, 128 x 128 x 512, failed due to its overwhelming computing requirements. The registration did not succeed for a resolution coarser than 64 x 64 x 64, allowing very little scope for multiresolution registration. In Table 4, we present results of registration at 64 x 64 x 64 and 64 x 64 x 256 resolution levels. As expected, the average distance error was smaller for more finely sampled data. A limited two-level registration approach, which will first perform a coarser registration at 64 x 64 x 64 resolution followed by the same at 64 x 64 x 256, reducing the convergence time, is possible.

The convergence time of registration is important for its clinical acceptability and suitability for specific applications. Although not optimized, execution times for the registration program written in Matlab (The Mathworks, Inc., Natick, MA) and running on a 450 MHz Pentium-II personal computer with 512 MB of memory are reported in Table 5.

The timings were four times as much for 64 x 64 x 256 resolution data sets. We expect the timing to improve by at least factor of two by porting Matlab code to compiled language (C/C++) code. An additional factor-of-two improvement will result by running the program on a late-model computer with a 900 MHz or higher CPU clock.

VI. DISCUSSION

We have demonstrated that it is feasible to register ultrasound volumes, despite their relatively poor image quality and the presence of several image artifacts, using the mutual information property of voxel similarity. Furthermore, we have demonstrated that the image registration can recover a global deformation as simple as a rigid-body transformation and as complex as an affine transformation. We discuss the features and performance of our algorithm and the likely clinical applications below.

A successful approach to ultrasound image registration should be fundamentally 3D and compensate for both rigid and nonrigid deformations of the underlying organ. Moreover, it should be robust, accurate and fast and should require minimal user intervention. Our method is 3D and accommodates global nonrigid deformation of an organ between image acquisitions. We note that the shape mismatch between two representations of an organ from two different instants could arise from two sources. The first is the deformation due to patient positioning and the other, any physiological and pathological changes in the organ itself. For successful temporal comparison, it is imperative that the registration account for only the patient positioning component, not the physiological and pathological changes. The shape of a deformable soft tissue organ cannot be the same between differing patient positions such as the supine and the prone due to gravity and pressure changes from neighboring organs. Even small variations in a given position (supine, for example) between image acquisitions could contribute to deformation. Although this deformation will in general be nonrigid, it is expected to be global primarily because gravity, the major contributor, is uniform. The physiological and pathological changes such as tissue growth or decay, on the other hand, are expected to be local. Although a totally elastic registration is possible within the mutual information-based image registration framework, a caveat with elastic image registration is the loss of tissue changes that are of clinical interest. Constraining deformation to a global transformation during registration, at least in principle, allows serial follow-up without patient positioning errors. An enhancement would exclude any local regions of pathology from the registration process. Meyer *et al.* [19] showed that the affine transformation was adequate for registering images of the breast, a highly deformable organ. A totally elastic registration, however, does make sense in intermodality registration, where local tissue changes are not an issue and the goal is fusion of complementary information.

Our work is anticipatory. Image registration is one of the core image analysis technologies we are developing in parallel with our development of a high-resolution real-time 3D ultrasound scanner based on synthetic aperture beam forming. The utilization of Volumetrics data sets was to enable registration algorithm development, realizing that the accuracy and the capture range of the algorithm will improve with higher resolution images produced by our future scanner. We picture the accuracy and capture range derived from Volumetrics images as the lower bounds of the respective quantities. To estimate the expected accuracy better at the present time, we employed semi-synthetic volumes created from clinical images; however, parameters associated with the synthetic axis could not be determined very accurately. The accuracy (average distance error) by excluding parameters associated with synthetic axis was found to be 0.86 mm in the rigid-body case and 1.6 mm in the full affine case. These numbers serve as the upper bounds of

accuracy and are comparable to the voxel size. Given the poor image quality of ultrasound, achievement of subvoxel accuracy in registration is not anticipated.

In general, and as discussed by Carrillo *et al.* [16] also, the accuracy of image registration is difficult to assess for organs that can deform or move with respect to the exterior of the body, which includes virtually all organs imaged by ultrasound. An external marker-based approach that has been used successfully in the validation of brain image registration does not apply to the case of deformable organs, thus limiting one to comparing internal landmarks, concurrence with experts and recovery of user-introduced transformations. We took the latter approach, registering two frames closely spaced in time. The temporal proximity allowed us to assume equality of anatomy without co-related noise. Despite real-time acquisition of Volumetrics images and regular heartbeat in the case of semi-synthetic data, a slight difference in anatomy, which might have affected the reported accuracy, could not be ruled out. Overall, the accuracy was acceptable in the RB and RB + US transformation modes. Studholme *et al.* [15] reported 3 mm and 4 degrees (7 mm displacement 100 mm away from the center of axis) as the limits of accuracy when human experts performed a manual registration of MRI and Positron Emission Tomography (PET) brain images. The specified numbers for translational accuracy were comparable to the voxel size of PET, the lower resolution image in the pair. Although no such data are available for ultrasound images of the heart, optimistic limits of human accuracy could be considered equivalent to the voxel size. Using the voxel dimension yardstick, the accuracy is satisfactory in RB and RB + US transformation modes. The same cannot be claimed for RB + NS and AT modes. We do point out that all the accuracy numbers reported by us included the effect of outliers as well. A quantitative measure to separate cases that converged to an acceptable solution from those that did not was difficult to define. The accuracy numbers, therefore, also reflect the robustness of registration. The RB + NS and AT modes are thus less robust. The starting misalignment was kept at 6 parameter units for all experiments. We also note that the capture range was estimated to be 6 parameters units for affine transformation. The lack of robustness in the affine case may in part be due to the starting misalignment bordering the limit of convergence.

We have found simplex optimization to be most successful with ultrasound data even though Powell's method has typically been used in voxel similarity registration in reported studies [15, 16]. Simplex optimization is, however, time-consuming as reflected in the execution times we have reported. Since the use of voxel similarity registration for real-time 3D ultrasound images is not a widely investigated area, our immediate focus was effectiveness, not speed. Encouraged by our result and convinced that a generalized nonrigid registration of ultrasound volumes is indeed feasible, we are exploring ways to improve execution of the algorithm. We have looked at the possibility of well-known multiresolution strategy; however, the existing image quality of ultrasound does not allow effective registration at lower resolutions. Since the algorithm spends 80% of the time in resampling the secondary, we are investigating hardware acceleration for a factor-of-five increase in speed. Since the execution time is proportional to the dimensionality of the parameter space, and the capture range shrinks with the complexity of transformation, a progressive refinement approach, which will both save time and improve accuracy, is desirable. Starting with a rigid-body transformation, one would proceed to affine transformation through intermediate limited affine transformation modes while the size of the search space progressively becomes narrower.

Clinical applications of a generalized, accurate and robust 3D image registration technique could be many. At the very least, it would allow comparison, in quantitative terms, of the images of the organs such as the kidney and the liver. Image quality could be improved by registration and subsequent superimposition of several successive scans. Many cardiac applications are possible as well. One specific application to which we are applying the developed techniques is the alignment of pre- and post-stress ultrasound 3D images of the heart. Once registered, a side-by-side presentation of pre- and post-stress images along any arbitrary orientation is possible, allowing a physician to perform accurate and comprehensive diagnosis.

VII. CONCLUSION

We have demonstrated that mutual information-based registration, previously developed for multimodality registration of brain images, is effective for registration of 3D ultrasound images. Furthermore, given the deformable nature of any organ typically imaged with ultrasound, we can apply the same framework for nonrigid deformations such as rigid-body transformation with both uniform and nonuniform scaling and affine transformation. The accuracy of the registration is comparable to voxel dimensions; however, the capture range shrinks with the complexity of deformation. We expect both accuracy and capture range to improve with higher resolution images to be produced by a synthetic aperture real-time 3D ultrasound scanner being developed at our institution.

APPENDIX

TRANSFORMATION MATRIX FORMULATION

A generalized affine transformation is the product of scaling (S), shearing (H), rotation (R) and translation (D) matrices. For the translation vector $\{d_x, d_y, d_z\}$ and the scaling vector $\{s_x, s_y, s_z\}$, the translation and scaling matrices are expressed as

$$D = \begin{bmatrix} 1 & 0 & 0 & d_x \\ 0 & 1 & 0 & d_y \\ 0 & 0 & 1 & d_z \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad S = \begin{bmatrix} s_x & 0 & 0 & 0 \\ 0 & s_y & 0 & 0 \\ 0 & 0 & s_z & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The rotation matrix is the product of three matrices representing individual rotation about the x, y and z axes, respectively, by angles ϕ_x , ϕ_y and ϕ_z . $R = R_z * R_y * R_x$, where

$$R_z = \begin{bmatrix} \cos \phi_z & -\sin \phi_z & 0 & 0 \\ \sin \phi_z & \cos \phi_z & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad R_y = \begin{bmatrix} \cos \phi_y & 0 & -\sin \phi_y & 0 \\ 0 & 1 & 0 & 0 \\ \sin \phi_y & 0 & \cos \phi_y & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad \text{and}$$

$$R_x = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \phi_x & -\sin \phi_x & 0 \\ 0 & \sin \phi_x & \cos \phi_x & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

The shearing matrix is a product of six matrices of the form

$$H_{ab} = \begin{bmatrix} 1 & h_{xy} & h_{xz} & 0 \\ h_{yx} & 1 & h_{yz} & 0 \\ h_{zx} & h_{zy} & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

where $h_{cd} = \begin{cases} -\tan \theta_{ab}, & \text{if } a=c \text{ and } b=d \\ 0, & \text{otherwise} \end{cases}$, a, b, c and d assume values x, y and z .

Each pair of axes produces two shears, θ_{ab} and θ_{ba} , but only one of them is unique. Consequently, only three out of six shear parameters need to be used for transformation matrix formulation for registration. To prove this, let us consider the transformation H_3 achieved with the set of three redundant shears.

$$H_3 = H_{yx} * H_{xz} * H_{zy} = \begin{bmatrix} 1 & 0 & -\tan \theta_{xz} & 0 \\ -\tan \theta_{yx} & 1 & 0 & 0 \\ \tan \theta_{yx} * \tan \theta_{zy} & -\tan \theta_{zy} & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}.$$

Let us consider a second transformation H_9 incorporating scaling, three rotations and the other three shear matrices.

$$H_9 = S * H_{xy} * H_{zx} * H_{yz} * R_x * R_y * R_z =$$

$$= \begin{bmatrix} s_x(\cos \phi_y - (\sin \phi_x \tan \theta_{xy} + \cos \phi_x \tan \theta_{xy} \tan \theta_{yz}) \sin \phi_y) \cos \phi_z - s_x(\cos \phi_x \tan \theta_{xy} - \sin \phi_x \tan \theta_{xy} \tan \theta_{yz}) \sin \phi_z & \dots \\ s_y(\sin \phi_x + \cos \phi_x \tan \theta_{yz}) \sin \phi_y \cos \phi_z + s_y(\cos \phi_x - \sin \phi_x \tan \theta_{yz}) \sin \phi_z & \dots \\ -s_z(\cos \phi_y \tan \theta_{zx} - \cos \phi_x \sin \phi_y) \cos \phi_z + s_z \sin \phi_x \sin \phi_z & \dots \\ 0 & \dots \\ \dots & \dots \\ -s_x(\cos \phi_y - (\sin \phi_x \tan \theta_{xy} + \cos \phi_x \tan \theta_{xy} \tan \theta_{yz}) \sin \phi_y) \sin \phi_z - s_x(\cos \phi_x \tan \theta_{xy} - \sin \phi_x \tan \theta_{xy} \tan \theta_{yz}) \cos \phi_z & \dots \\ \dots & \dots \\ -s_y(\sin \phi_x + \cos \phi_x \tan \theta_{yz}) \sin \phi_y \sin \phi_z + s_y(\cos \phi_x - \sin \phi_x \tan \theta_{yz}) \cos \phi_z & \dots \\ \dots & \dots \\ s_z(\cos \phi_y \tan \theta_{zx} - \cos \phi_x \sin \phi_y) \sin \phi_z + s_z \sin \phi_x \cos \phi_z & \dots \\ 0 & \dots \\ \dots & \dots \\ s_x \sin \phi_y + s_x(\sin \phi_x \tan \theta_{xy} + \cos \phi_x \tan \theta_{xy} \tan \theta_{yz}) \cos \phi_y & 0 \\ \dots & \dots \\ -s_y(\sin \phi_x + \cos \phi_x \tan \theta_{yz}) \cos \phi_y & 0 \\ \dots & \dots \\ -s_z(\sin \phi_y \tan \theta_{zx} - \cos \phi_x \cos \phi_y) & 0 \\ \dots & \dots \\ 0 & 1 \end{bmatrix}$$

Element by element comparison of matrices H_3 and H_9 provides the following set of equations and solution.

$$\begin{aligned} H_9(2,3) &= 0 & \Rightarrow & \theta_{yz} = -\phi_x; \\ H_9(1,2) &= 0 & \Rightarrow & \theta_{xy} = \arctan(-\tan \phi_z \cos \phi_y \cos \phi_x); \\ H_9(2,1) * H_9(3,2) &= H_9(3,1) & \Rightarrow & \theta_{zx} = \arctan(\tan \phi_y \cos \phi_x); \\ H_9(1,1) &= 1 & \Rightarrow & s_x = \cos \phi_z / \cos \phi_y; \end{aligned}$$

$$\begin{array}{ll}
H_9(2,2) = 1 & \Rightarrow s_y = \cos\phi_x / \cos\phi_z; \\
H_9(3,3) = 1 & \Rightarrow s_z = \cos\phi_y / \cos\phi_x. \\
H_9(2,1) = -\tan\theta_{yx} & \Rightarrow \phi_z = -\theta_{yx}; \\
H_9(1,3) = -\tan\theta_{xz} & \Rightarrow \phi_y = \arctan(\tan\theta_{xz} / \cos\phi_z); \\
H_9(3,2) = -\tan\theta_{zy} & \Rightarrow \phi_x = \arctan(-\tan\theta_{zy} / (\cos\phi_y \cos\phi_z)).
\end{array}$$

To summarize, three shearing parameters of H_3 can be expressed in terms of the other three shearing, three rotation and three scaling parameters.

ACKNOWLEDGMENTS

The authors thank Drs. James Thomas, Takahiro Shiota and Neil Greenberg of the Department of Cardiology at The Cleveland Clinic Foundation for making available the images used in the study, and Ms. Christine Kassuba of the Department of Biomedical Engineering for editorial assistance with the preparation of this manuscript.

REFERENCES

- [1] J. B. Maintz and M. A. Viergever, "A survey of medical image registration," *Medical Image Analysis*, vol. 2, pp. 1-36, 1998.
- [2] J. C. Gee, "On matching brain volumes," *Pattern Recognition*, vol. 32, pp. 99-111, 1999.
- [3] P. A. Freeborough, R. P. Woods, and N. C. Fox, "Accurate registration of serial 3D MR brain images and its application to visualizing change in neurodegenerative disorders," *Journal of Computer Assisted Tomography*, vol. 20, pp. 1012-22, 1996.
- [4] M. I. Miga, K. D. Paulsen, P. J. Hoopes, F. E. J. Kennedy, A. Hartov, and D. W. Roberts, "In vivo quantification of a homogeneous brain deformation model for updating preoperative images during surgery," *IEEE Transactions on Biomedical Engineering*, vol. 47, pp. 266-273, 2000.
- [5] J. L. Andersson, A. Sundin, and S. Valind, "A method for coregistration of PET and MR brain images," *Journal of Nuclear Medicine*, vol. 36, pp. 1307-15, 1995.
- [6] A. Fenster and D. B. Downey, "3-D ultrasound imaging: A review," *IEEE Engineering in Medicine and Biology*, vol. 15, pp. 41-51, 1996.
- [7] S. Berg, H. Torp, D. Martens, E. Steen, S. Samstad, I. Hoivik, and B. Olstad, "Dynamic three-dimensional freehand echocardiography using raw digital ultrasound data," *Ultrasound in Medicine and Biology*, vol. 25, pp. 745-753, 1999.
- [8] O. T. Von Ramm, S. W. Smith, and B. A. Carroll, "Advanced real-time volumetric ultrasound scanning," *J Ultrasound Med*, vol. 14, pp. S35, March 1995.
- [9] C. R. Hazard and G. R. Lockwood, "Theoretical assessment of a synthetic aperture beamformer for real-time 3-D imaging," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 46, pp. 972-980, 1999.
- [10] G. R. Lockwood, J. R. Talman, and S. S. Brunke, "Real-time 3-D ultrasound imaging using sparse synthetic aperture beamforming," *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, vol. 45, pp. 980-988, 1998.
- [11] T. Peters, B. Davey, P. Munger, R. Comeau, A. Evans, and A. Olivier, "Three-dimensional multimodal image-guidance for neurosurgery," *IEEE Transactions on Medical Imaging*, vol. 15, pp. 121-128, 1996.
- [12] R. T. Malison, E. G. Miller, R. Greene, G. McCarthy, D. S. Charney, and R. B. Innis, "Computer-assisted coregistration of multislice SPECT and MR brain images by fixed external fiducials," *Journal of Computer Assisted Tomography*, vol. 17, pp. 952-60, 1993.
- [13] F. Maes, A. Collignon, D. Vandermeulen, G. Marchal, and P. Suetens, "Multimodality image registration by maximization of mutual information," *IEEE Trans Med Imaging*, vol. 16, pp. 187-98, 1997.
- [14] J. West, J. M. Fitzpatrick, M. Y. Wang, B. M. Dawant, C. R. J. Maurer, R. M. Kessler, and R. J. Maciunas, "Retrospective intermodality registration techniques for images of the head: Surface-based versus volume-based," *IEEE Transactions on Medical Imaging*, vol. 18, pp. 144-150, 1999.
- [15] C. Studholme, D. L. Hill, and D. J. Hawkes, "Automated three-dimensional registration of magnetic resonance and positron emission tomography brain images by multiresolution optimization of voxel similarity measures," *Medical Physics*, vol. 24, pp. 25-35, 1997.
- [16] A. Carrillo, J. L. Duerk, J. S. Lewin, and D. L. Wilson, "Semiautomatic 3-D image registration as applied to interventional MRI liver cancer treatment," *IEEE Transactions on Medical Imaging*, vol. 19, pp. 175-185, 2000.
- [17] V. Zagrodsky, R. Shekhar, and J. F. Cornhill, "Mutual information based registration of cardiac ultrasound volumes," *Proceedings of SPIE - The International Society for Optical Engineering*, vol. 3979, pp. 1605-1614, 2000.
- [18] R. N. Rohling, A. H. Gee, and L. Berman, "Automatic registration of 3-D ultrasound images," *Ultrasound in Medicine and Biology*, vol. 24, pp. 841-854, 1998.
- [19] C. R. Meyer, J. L. Boes, B. Kim, P. H. Bland, G. L. Lecarpentier, J. B. Fowlkes, M. A. Roubidoux, and P. L. Carson, "Semiautomatic registration of volumetric ultrasound scans," *Ultrasound in Medicine and Biology*, vol. 25, pp. 339-347, 1999.
- [20] V. Zagrodsky, R. Shekhar, and J. F. Cornhill, "Multi-function extension of simplex optimization method for mutual information based registration of ultrasound volumes," accepted for presentation at SPIE's International Symposium on Medical Imaging, 2001.
- [21] J. A. Nelder and R. Mead, "A simplex method for function minimization," *J. Computer*, vol. 7, pp. 308-313, 1965.

TABLE 1

SUMMARY OF REGISTRATION RESULTS FOR ALL FIVE VOLUMETRICS DATA SETS.

Transformation Mode	Average Distance Error (mm)	rms Estimation Error of Transformation Parameter											
		d_x (mm)	d_y (mm)	d_z (mm)	ϕ_x (deg)	ϕ_y (deg)	ϕ_z (deg)	s_x (%)	s_y (%)	s_z (%)	θ_{xy} (deg)	θ_{yz} (deg)	θ_{zx} (deg)
RB	1.4	0.34	0.39	0.16	0.24	0.18	0.73						
		0.41	0.47	0.63	0.31	0.21	0.74						
RB+US	2.5	0.60	0.36	0.18	0.36	0.27	1.6	0.38					
		0.83	0.53	0.68	0.40	0.29	1.7	1.1					
RB+NS	6.2	1.5	0.79	2.3	1.1	1.6	1.3	5.9	2.6	7.2			
		1.6	0.96	2.4	1.2	1.7	1.4	6.2	3.3	7.2			
AT	14.1	3.2	2.5	4.0	2.7	5.3	6.2	7.3	5.1	15.0	5.4	5.0	4.5
		3.4	2.9	4.6	2.8	5.3	6.3	7.8	6.7	15.0	5.5	5.3	5.3

TABLE 2

SUMMARY OF REGISTRATION RESULTS FOR ALL SEMI-SYNTHETIC DATA SETS.

Transformation Mode	Average Distance Error (mm)	rms Estimation Error of Transformation Parameter											
		d_x (mm)	d_y (mm)	d_z (mm)	ϕ_x (deg)	ϕ_y (deg)	ϕ_z (deg)	s_x (%)	s_y (%)	s_z (%)	θ_{xy} (deg)	θ_{yz} (deg)	θ_{zx} (deg)
RB	2.0	0.17 0.33	0.38 0.99	0.25 0.27	0.52 0.62	0.13 0.13	0.35 1.2						
RB+US	1.7	0.18 0.38	0.56 0.68	0.26 0.35	0.22 0.58	0.18 0.18	0.42 0.93	0.35 0.35					
RB+NS	12.7	0.38 0.71	2.0 2.4	0.73 0.75	2.1 2.2	0.29 0.40	1.3 1.5	1.1 1.2	26.0 25.0	2.3 2.3			
AT	15.4	0.73 0.97	2.5 3.0	0.98 1.00	5.3 6.0	2.2 2.3	3.4 4.6	1.0 1.1	24.0 25.0	10.9 11.0	4.3 5.1	5.4 5.8	2.1 2.1

TABLE 3

SUMMARY OF REGISTRATION RESULTS FOR SEMI-SYNTHETIC DATA SETS
EXCLUDING TRANSFORMATION PARAMETERS AFFECTING DATA ALONG THE
SYNTHETIC (Y) AXIS.

Transformation Mode	Average Distance Error (mm)	rms Estimation Error of Transformation Parameter											
		d_x (mm)	d_y (mm)	d_z (mm)	ϕ_x (deg)	ϕ_y (deg)	ϕ_z (deg)	s_x (%)	s_y (%)	s_z (%)	θ_{xy} (deg)	θ_{yz} (deg)	θ_{zx} (deg)
RB	0.86	0.40 0.41		0.47 0.52		0.20 0.20							
RB+US													
RB+NS	1.3	0.48 0.51		0.53 0.61		0.30 0.31		0.37 0.37		0.43 0.48			
AT	1.6	0.66 0.68		0.65 0.80		0.25 0.25		0.60 0.61		1.10 1.20			0.39 0.40

TABLE 4

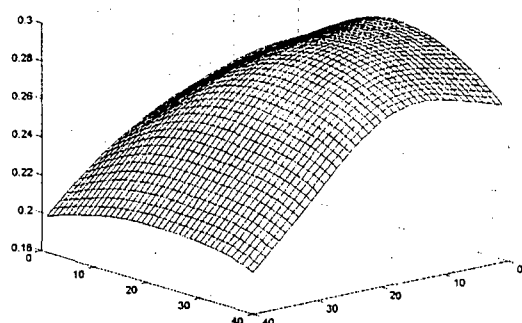
ACCURACY OF REGISTRATION FOR VOLUMETRICS DATA SETS AT TWO LEVELS
OF RESOLUTION.

Data Resolution	Deformation mode	Average Distance Error (mm)
64 x 64 x 64	RB	2.0
	RB + US	4.1
	RB + NS	8.2
	AT	14.2
64 x 64 x 256	RB	1.6
	RB + US	2.7
	RB + NS	6.2
	AT	11.6

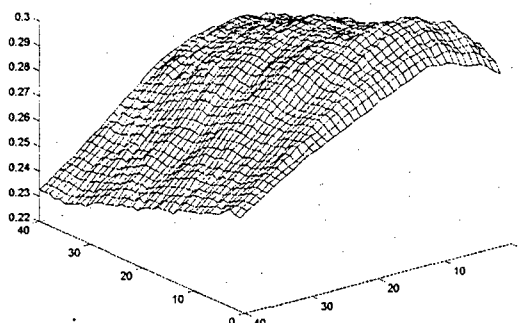
TABLE 5

EXECUTION TIME FOR 64 X 64 X 64 RESOLUTION VOLUMETRICS DATA.

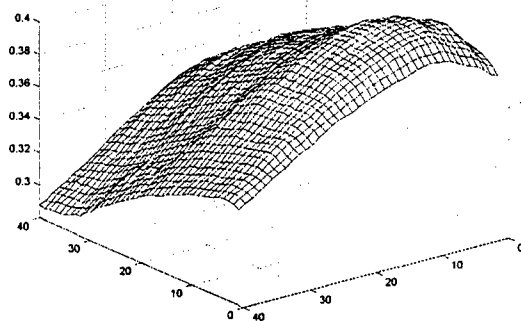
Transformation mode	Execution time (min)
RB	6 - 8
RB + US	8 - 10
RB + NS	12 - 15
AT	20 - 25



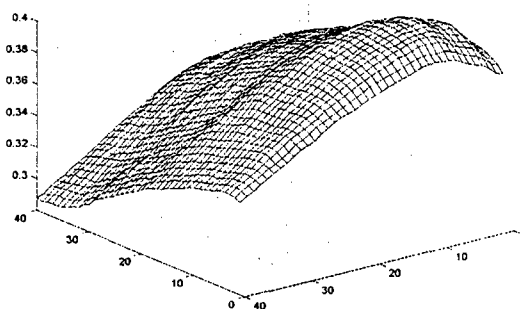
(a) MRI + SPECT



(b) US + US



(c) US + US following medial filtering



(d) US + US, median filtering + 6 bit quantization

Figure 1 Mutual information as a function of misalignment for an MRI and SPECT image pair (panel (a)) and a 3D ultrasound (US) image pair (panel (b)) with no preprocessing. Panels (c) and (d) show the mutual information surface plots, following medial filtering, for 8 and 6 bits of intensity quantization, respectively. Note the inherent smoothness of the mutual information surface for the non-ultrasound image pair in panel (a) and the roughness of the same for the ultrasound image pair in panel (b). The effect of preprocessing to smooth the surface through median filtering and then intensity quantization is apparent in panels (c) and (d).

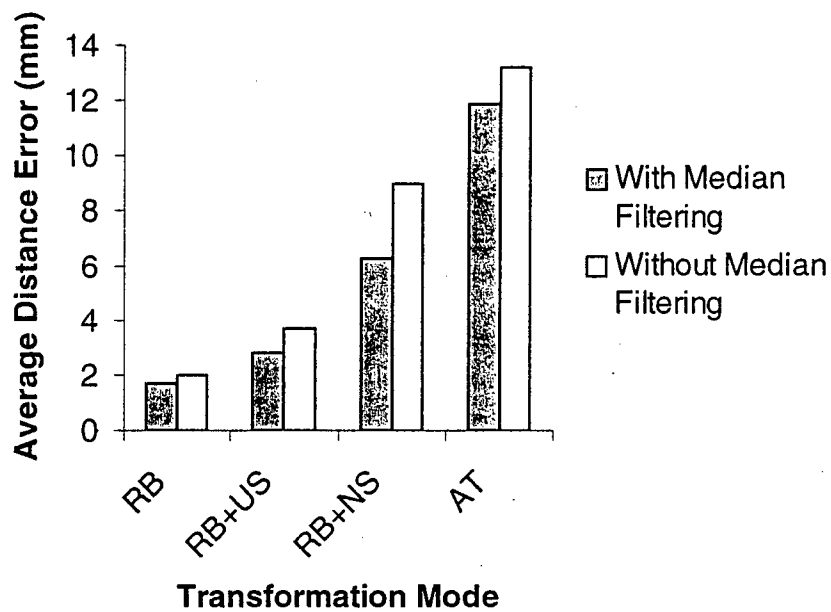


Figure 2 Average distance error with and without median filtering for the four transformation modes of image registration. Note the error is smaller with median filtering in each of the four transformation modes.

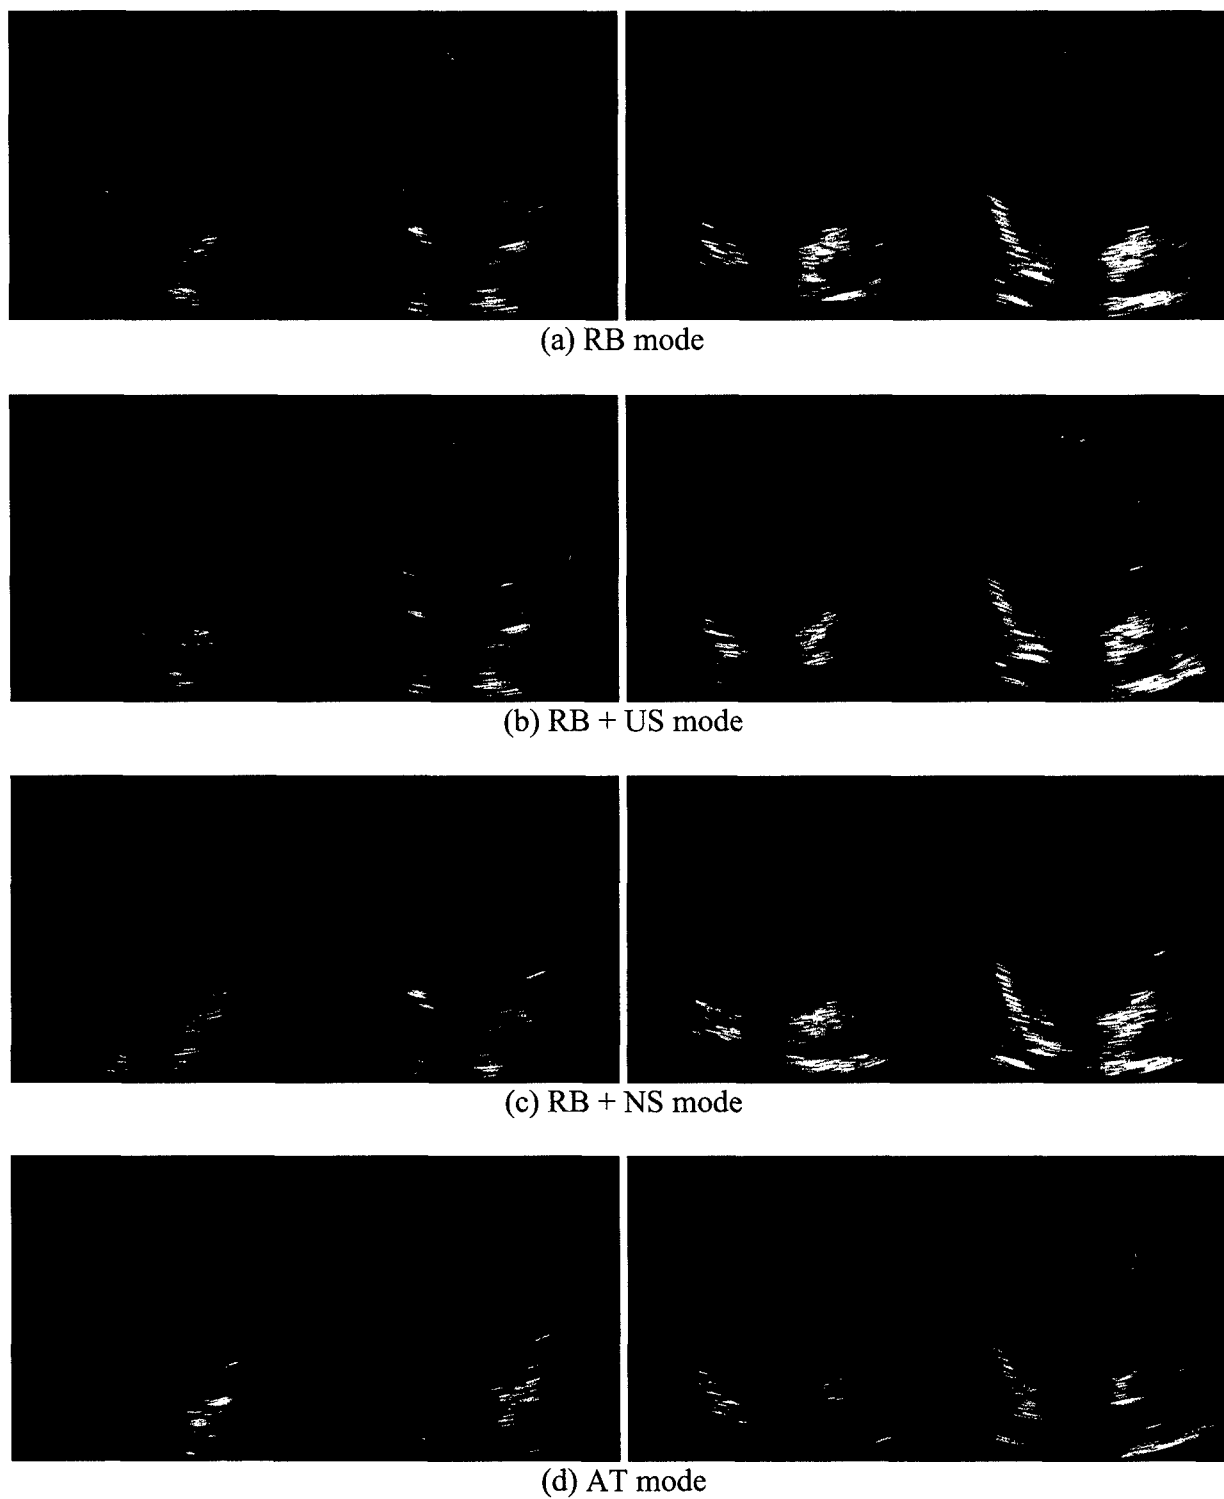


Figure 3 Fused primary and secondary Volumetrics data sets before and after registration for all four deformation modes. Each row has a pair of fused images before and after registration for two orthogonal cross-sections of the 3D data.

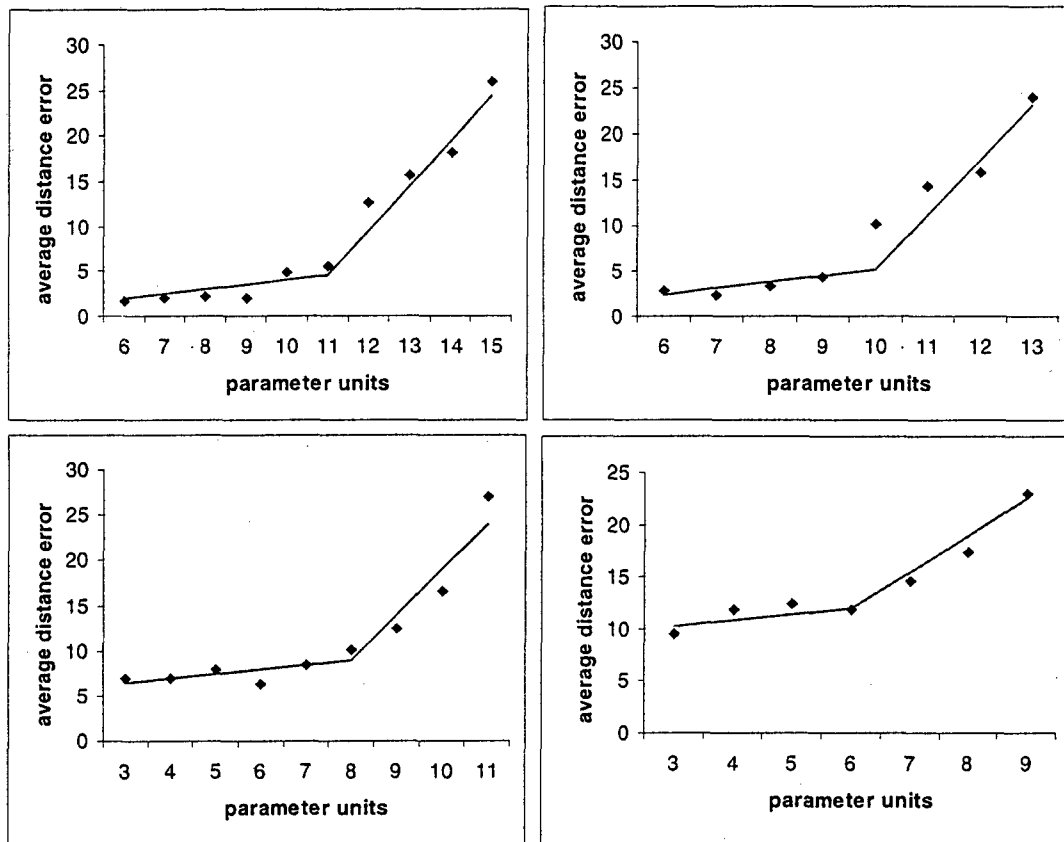
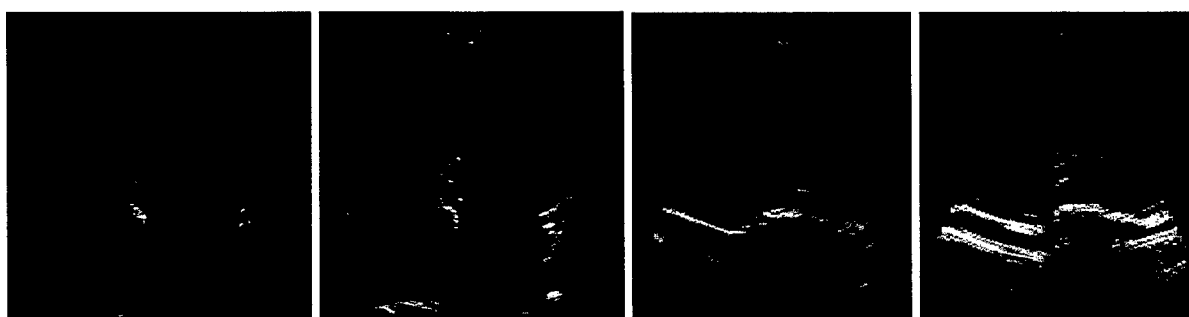
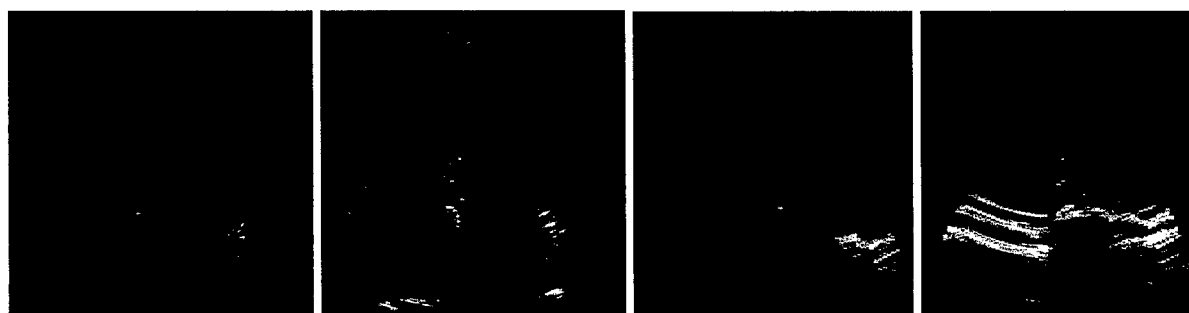


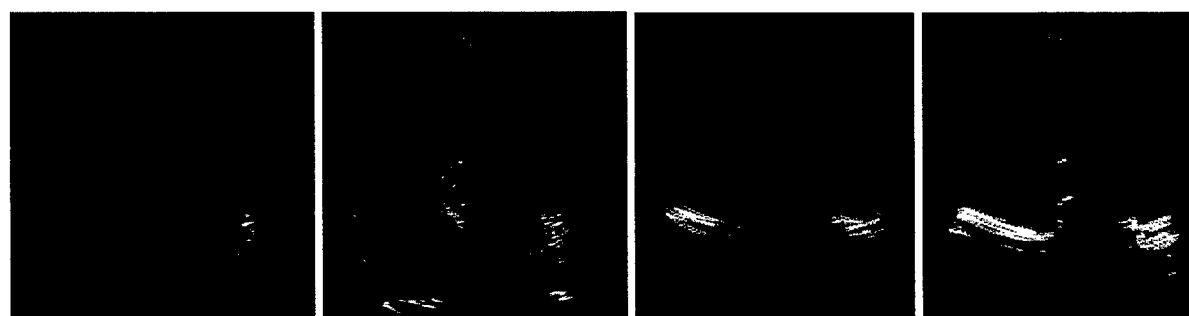
Figure 4 Accuracy of registration as a function of misalignment. The panels (a), (b), (c) and (d) correspond to RB, US, NUS and AT transformations, respectively. Note that the capture range, defined by the point of inflexion, shrinks with increasing complexity of transformation.



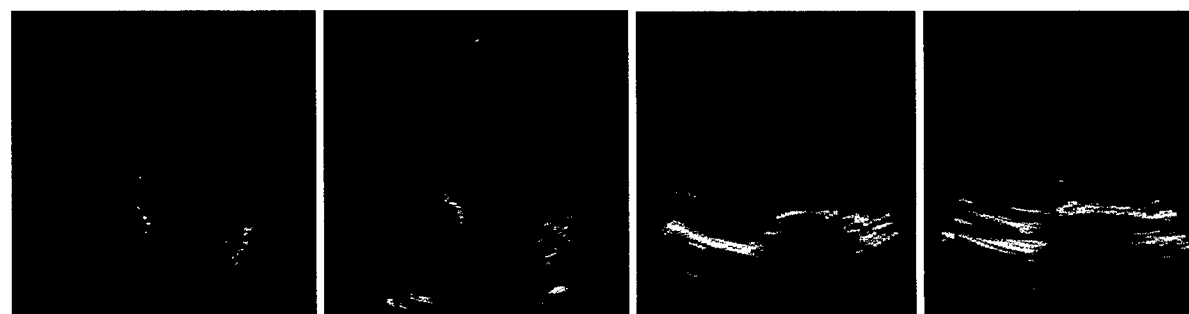
(a) RB mode



(b) RB + US mode



(c) RB + NS mode



(d) AT mode

Figure 5 Fused primary and secondary semi-synthetic data before and after registration for all four transformation modes. Each row shows a pair of fused images before and after registration for two orthogonal cross-sections.

Appendix 4: Abstract of the paper to be presented at SPIE Medical Imaging Symposium in San Diego in February 2001 - Vladimir Zagrotsky, Raj Shekhar, and J. Fredrick Cornhill, "Multifunction extension of simplex optimization method for mutual information based registration of ultrasound volumes."

(To be presented at SPIE Medical Imaging Symposium in San Diego, February 2001)

Multi-Function Extension of Simplex Optimization Method for Mutual Information Based Registration of Ultrasound Volumes

Vladimir A. Zagrodsky, Raj Shekhar, J. Fredrick Cornhill (The Cleveland Clinic Foundation, Cleveland, OH 44195)

Mutual information has been demonstrated to be an accurate and reliable criterion function to perform registration of medical data. Due to speckle noise, ultrasound volumes do not provide a smooth mutual information function. Consequently the optimization technique used must be robust to avoid local maxima and converge on the desired global maximum eventually. While the well-known simplex direct search optimization uses a single criterion function, our extension to multi-function optimization uses three criterion functions, namely mutual information computed at three levels of intensity quantization and hence three degrees of speckle noise suppression. Registration was performed with rigid as well as simple non-rigid transformation modes for real-time 3D ultrasound datasets of the left ventricle. Pairs of frames corresponding to the most stationary end-diastolic cardiac phase were chosen, and the initial misalignment was introduced artificially. The multi-function simplex optimization reduced the failure rate by a factor of two in comparison to the standard direct search, while the average accuracy for the successful cases was unchanged. A more robust registration resulted from the parallel use of criterion functions. The additional computational cost was negligible, as each of the three implementations of mutual information used the same joint histogram and required no extra spatial transformation.



DEPARTMENT OF THE ARMY
US ARMY MEDICAL RESEARCH AND MATERIEL COMMAND
504 SCOTT STREET
FORT DETRICK, MARYLAND 21702-5012

REPLY TO
ATTENTION OF:

MCMR-RMI-S (70-1y)

28 July 03

MEMORANDUM FOR Administrator, Defense Technical Information
Center (DTIC-OCA), 8725 John J. Kingman Road, Fort Belvoir,
VA 22060-6218


SUBJECT: Request Change in Distribution Statement

1. The U.S. Army Medical Research and Materiel Command has reexamined the need for the limitation assigned to technical reports written for this Command. Request the limited distribution statement for the enclosed accession numbers be changed to "Approved for public release; distribution unlimited." These reports should be released to the National Technical Information Service.

2. Point of contact for this request is Ms. Kristin Morrow at DSN 343-7327 or by e-mail at Kristin.Morrow@det.amedd.army.mil.

FOR THE COMMANDER:

Encl


PHYLLIS M. RINEHART
Deputy Chief of Staff for
Information Management

ADB233865	ADB264750
ADB265530	ADB282776
ADB244706	ADB286264
ADB285843	ADB260563
ADB240902	ADB277918
ADB264038	ADB286365
ADB285885	ADB275327
ADB274458	ADB286736
ADB285735	ADB286137
ADB286597	ADB286146
ADB285707	ADB286100
ADB274521	ADB286266
ADB259955	ADB286308
ADB274793	ADB285832
ADB285914	
ADB260288	
ADB254419	
ADB282347	
ADB286860	
ADB262052	
ADB286348	
ADB264839	
ADB275123	
ADB286590	
ADB264002	
ADB281670	
ADB281622	
ADB263720	
ADB285876	
ADB262660	
ADB282191	
ADB283518	
ADB285797	
ADB269339	
ADB264584	
ADB282777	
ADB286185	
ADB262261	
ADB282896	
ADB286247	
ADB286127	
ADB274629	
ADB284370	
ADB264652	
ADB281790	
ADB286578	